LSE Research Online

Christian List

Distributed cognition: a perspective from social choice theory

Book section

http://eprints.lse.ac.uk

Distributed Cognition: A Perspective from Social Choice Theory

by

CHRISTIAN LIST[1]

Distributed cognition refers to processes which are (i) cognitive and (ii) distributed across multiple agents or devices rather than performed by a single agent. Distributed cognition has attracted interest in several fields ranging from sociology and law to computer science and the philosophy of science. In this paper, I discuss distributed cognition from a social-choice-theoretic perspective. Drawing on models of judgment aggregation, I address two questions. First, how can we model a group of individuals as a distributed cognitive system? Second, can a group acting as a distributed cognitive system be 'rational' and 'track the truth' in the outputs it produces? I argue that a group's performance as a distributed cognitive system depends on its 'aggregation procedure' – its mechanism for aggregating the group members' inputs into collective outputs – and I investigate the properties of an aggregation procedure that matter.

1 Introduction

'Distributed cognition' refers to processes with two properties. First, they are *cognitive*, i.e. they involve forming certain representations of the world. Second, they are not performed by a single (human) agent, but are *distributed* across multiple (human) agents or (technical) devices. Distributed cognition has attracted interest in several fields, ranging from law (e.g. jury decision making) and sociology (e.g. information processing in organizations) to computer science (e.g. GRID computing) and the philosophy of science (e.g. expert panels).

An influential account of distributed cognition is HUTCHINS's [1995] study of navigation on a US Navy ship. Hutchins describes the ship's navigation as a process of distributed cognition. It is a cognitive process in that it leads to representations of the ship's position and movements in its environment. It is distributed in that there is no single individual on the ship who performs the complex navigational task alone, but the task is performed through the interaction of many individuals, together with technical instruments. At any given time, no single individual may be fully aware of the navigational process in its entirety. Thus, on Hutchins's account, the ship's navigation is performed not at the level of a single individual – a 'chief navigator' – but at the level of a larger system.

In the philosophy of science, GIERE [2002] argues that many scientific practices, especially large-scale collaborative research practices, involve distributed cognition, as these practices are "situation[s] in which one or more individuals reach a cognitive outcome either by combining individual knowledge not initially shared with the others or by interacting with artefacts organized in an appropriate way (or both)" [2002, p. 641]. He distinguishes between 'distributed' and 'collective' cognition, where the first is more general than the second. Distributed cognition includes not only cases of collective cognition, where a cognitive task is distributed across multiple individuals, but also cases where such a task is distributed between a single

individual and an artifact, such as a technical instrument.[2] While researchers often compete with one another, collectively distributed cognition is a phenomenon associated with more cooperative practices within research groups or communities.

KNORR CETINA [1999] provides a case study of distributed cognition in science. Studying high-energy physics research at the European Center for Nuclear Research (CERN), she observes that experiments, which lead to cognitive outcomes, involve many researchers and technicians, using complex technical devices, with a substantial division of labour, expertise, and authority. She describes this research practice as "something like distributed cognition" [p. 25, cited in GIERE [2002]].[3]

Other instances of distributed cognition in science can be found in multi-member expert committees. For example, in 2000, the National Assessment Synthesis Team, an expert committee commissioned by the US Global Change Research Program with members from governments, universities, industry and non-governmental organizations, presented a report on climate change.[4] Such a committee's work is cognitive in that it involves the representation of certain facts about the world; and it is distributed in that it involves a division of labour between multiple committee members and a pooling of different expertise and judgments. Here it may be more plausible to ascribe authorship of the report to the committee as a whole rather than any particular committee member.

In this paper, I discuss collectively distributed cognition from the perspective of social choice theory. Social choice theory can provide a general theory of the aggregation of multiple (individual) inputs into single (collective) outputs, although it is usually applied to the aggregation of preferences. Drawing on social-choice-theoretic models from the emerging theory of judgment aggregation (e.g. LIST AND PETTIT [2002], [2004]; PAULY AND VAN HEES [2005]; DIETRICH [2005]; BOVENS AND RABINOWICZ [2005]; LIST [2005a], [2005b], [2006]), I address two questions. First, how can we model a group of individuals as a distributed cognitive system? And, second, can a group acting as a distributed cognitive system be rational and track the truth in its cognitive outputs?

I argue that a group's performance as a distributed cognitive system depends crucially on its organizational structure, and a key part of that organizational structure is the group's 'aggregation procedure', as defined in social choice theory. An 'aggregation procedure' is a mechanism a multi-member group can use to combine ('aggregate') the judgments or representations held by the individual group members into judgments or representations endorsed by the group as a whole. I investigate the ways in which a group's aggregation procedure affects its capacity to be rational and to track the truth in the outputs it produces as a distributed cognitive system.

---

[2] Collective cognition is "[a] special case of distributed cognition, in which two or more individuals reach a cognitive outcome simply by combining individual knowledge not initially shared with others" (GIERE [2002, p. 641]).

[3] Knorr Cetina also studies research practices in molecular biology, but argues that here research is more individualized than in high energy physics and "the person remains the epistemic subject" [p. 217, cited in GIERE [2002]]. GIERE [2002, especially p. 643] responds that, while there may be less *collective* cognition in molecular biology than in high energy physics, there may still be *distributed* cognition, "where the cognition is distributed between an individual person and an instrument".

[4] The title of the report is "Climate Change Impacts on the United States: The Potential Consequences of Climate Variability and Change". See http://www.usgcrp.gov/.

My discussion is structured as follows. I begin with some introductory remarks about modelling a group as a distributed cognitive system in section 2 and introduce the concept of an aggregation procedure in section 3. The core of my discussion consists of sections 4 and 5, in which I discuss a group's capacity to be rational and to track the truth in its cognitive outputs, respectively. In section 6, I draw some conclusions.

2 Modelling a group as a distributed cognitive system

When does it make sense to consider a group of individuals as a distributed cognitive system rather than a mere collection of individuals? First, the group must count as a well-demarcated *system*, and, second, it must count as a system that produces *cognitive outputs*.

The first condition is met if and only if the group's collective behaviour is sufficiently integrated. A well organized expert panel, a group of scientific collaborators or the monetary policy committee of a central bank, for example, may have this property, whereas a random crowd of people at London's Leicester Square lacks the required level of integration. And the second condition is met if and only if the group is capable of producing outputs that have representational content; let me call these outputs 'collective judgments'. If a group's organizational structure – e.g. its procedures for generating a joint report – allows the group to make certain joint declarations that count as collective judgments, then the group has this property, whereas a group without any formal or informal organization, such as a random crowd at Leicester Square, lacks the required capacity.

At first sight, we may be reluctant to attribute judgments to groups over and above their individual members. But, as GOLDMAN [2004, p. 12] has noted, in ordinary language, groups or collective organizations are often treated as subjects for the attribution of judgments. Goldman's example is the recent debate on what the FBI as a collective organization did or did not "know" prior to the terrorist attacks of 9/11. In addition to the literature on distributed cognition, there is now a growing literature in philosophy that considers conditions under which groups are sufficiently integrated to produce outputs that we normally associate with rational agency (e.g. ROVANE [1998]; PETTIT [2003]; LIST AND PETTIT [2005a], [2005b]). Roughly, a sufficient level of integration is given in those cases in which it is pragmatically and explanatorily useful to describe the group's outputs in intentional terms (DENNETT [1987]), namely as the group's 'beliefs', 'judgments', 'commitments' or 'knowledge'. Arguably, this condition is satisfied by those groups that Hutchins, Giere, Knorr-Cetina and others have described as distributed cognitive systems.

In short, a necessary condition for distributed cognition in a group is the presence of an organizational structure that allows the group to produce collective judgments, i.e. collective outputs with representational content. Once this necessary condition is met, the group's performance as a distributed cognitive system depends on the nature of that organizational structure.

Consequently, to construct a model of a group as a distributed cognitive system, we need to represent not only the individual group members, but also the group's organizational structure. In the next section, I illustrate how we can think about this organizational structure in terms of a simple social-choice-theoretic model.

3 The concept of an aggregation procedure

How can we think about a group's organizational structure? Let me introduce the concept of an 'aggregation procedure' to represent (a key part of) a group's organizational structure. As defined in the theory of judgment aggregation (LIST AND PETTIT [2002], [2004]; LIST [2006]), an aggregation procedure is a mechanism by which a group can generate collective judgments on the basis of the group members' individual judgments (illustrated in table 1). Formally, an aggregation procedure is a function which assigns to each combination of individual judgments across the group members a corresponding set of collective judgments. A simple example is 'majority voting', whereby a group judges a given proposition to be true whenever a majority of group members judges it to be true. Below I discuss several other aggregation procedures.

Table 1: An aggregation procedure

Input (individual judgments)

Aggregation procedure

Output (collective judgments)
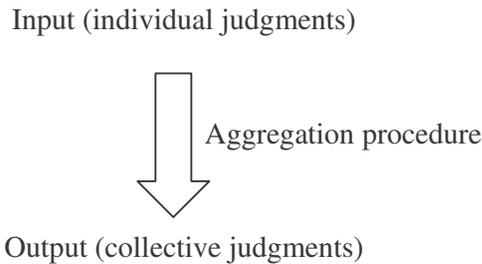
Of course, an aggregation procedure captures only part of a group's organizational structure (which may be quite complex), and there are also multiple ways (both formal and informal ones) in which a group might implement such a procedure. Nonetheless, as argued below, aggregation procedures are key factors in determining a group's performance as a distributed cognitive system.

In the next section, I ask what properties a group's aggregation procedure must have for the group to be rational as a distributed cognitive system – specifically, consistent, but also complete, in its collective judgments – and in the subsequent section, I ask what properties it must have for the group to track the truth in these judgments. Both discussions illustrate that a group's performance as a distributed cognitive system depends on its aggregation procedure.

4 Rationality in a distributed cognitive system

Suppose a group is given a cognitive task involving the formation of collective judgments on some propositions. Can the group ensure the consistency of these judgments?

4.1 A majoritarian inconsistency

Consider an expert committee that has to prepare a report on the health consequences of air pollution in a big city, especially pollution by particles smaller than 10 microns in diameter. This is an issue on which there has recently been much debate in Europe. The experts have to make judgments on the following propositions:

p: The average particle pollution level exceeds $50\mu gm^{-3}$ (micrograms per cubic meter air).

$p{\rightarrow}q$: If the average particle pollution level exceeds $50\mu gm^{-3}$, then residents have a significantly increased risk of respiratory disease.

q: Residents have a significantly increased risk of respiratory disease.

All three propositions are complex factual propositions on which the experts may disagree.[5] Suppose the group uses majority voting as its aggregation procedure, i.e. the collective judgment on each proposition is the majority judgment on that proposition, as defined above. Now suppose the experts' individual judgments are as shown in table 2.

Table 2: A majoritarian inconsistency

	p	*p→q*	*q*
Individual 1	True	True	True
Individual 2	True	False	False
Individual 3	False	True	False
Majority	True	True	False

Then a majority of experts judges *p* to be true, a majority judges *p→q* to be true, and yet a majority judges *q* to be false, an inconsistent collective set of judgments. The expert committee fails to be rational in the collective judgments it produces as a distributed cognitive system.

This problem – sometimes called a 'discursive dilemma' – illustrates that, under the initially plausible aggregation procedure of majority voting, a group acting as a distributed cognitive system may not achieve consistent collective judgments even when all group members hold individually consistent judgments (PETTIT [2001]; LIST AND PETTIT [2002], [2004]; LIST [2006]; the problem originally goes back to the so-called 'doctrinal paradox' first identified by Kornhauser and Sager [1986]).

Is the present example just an isolated artefact, or can we learn something more general from it?

4.2 An impossibility theorem

Consider again any group of two or more individuals that is given the cognitive task to form collective judgments on a set of non-trivially interconnected propositions, as in the expert committee example.[6] Call an agent's judgments on these propositions 'complete' if, for each proposition-negation pair, the agent judges either the proposition or its negation to be true; and call these judgments 'consistent' if the set of propositions judged to be true by the agent is a consistent set in the standard sense of propositional logic.[7]

---

[5] Propositions *p* and *p→q* can be seen as 'premises' for the 'conclusion' *q*. Determining whether *p* is true requires an evaluation of air quality measurements; determining whether *p→q* is true requires an understanding of causal processes in human physiology; finally, determining whether *q* is true requires a combination of the judgments on *p* and *p→q*.

[6] Following LIST [2006], a set of propositions is 'non-trivially interconnected' if it is of one of the following forms (or a superset thereof): (i) it includes *k*>1 propositions p_1, ..., p_k and either their conjunction 'p_1 and ... and p_k' or their disjunction 'p_1 or p_2 or … or p_k' or both (and the negations of all these propositions); (ii) it includes *k*>1 propositions p_1, ..., p_k, another proposition *q* and either the proposition '*q* if and only if (p_1 and ... and p_k)' or the proposition '*q* if and only if (p_1 or p_2 or … or p_k)' or both (and negations); (iii) it includes propositions *p*, *q* and *p→q* (and negations).

[7] This consistency notion is stronger than that in LIST AND PETTIT [2002]. But when the present consistency notion is used, no deductive closure requirement needs to be added.

Suppose now that each individual holds complete and consistent judgments on these propositions, and that the collective judgments are also required to be complete and consistent. One can then prove the following impossibility result (for a discussion of parallels and disanalogies between this result and ARROW's (1951) classical theorem, see LIST AND PETTIT [2004] and DIETRICH AND LIST [2005a]).

Theorem (LIST AND PETTIT [2002]). There exists no aggregation procedure generating complete and consistent collective judgments that satisfies the following three conditions simultaneously:

Universal domain. The procedure accepts as admissible input any logically possible combinations of complete and consistent individual judgments on the propositions.

Anonymity. The judgments of all individuals have equal weight in determining the collective judgments.

Systematicity. The collective judgment on each proposition depends only on the individual judgments on that proposition, and the same pattern of dependence holds for all propositions.

In short, majority voting is not the only aggregation procedure that runs into problems like the one illustrated in table 2 above. Any procedure satisfying universal domain, anonymity and systematicity does so. If these conditions are regarded as indispensable requirements on an aggregation procedure, then one has to conclude that a multi-member group acting as a distributed cognitive system cannot ensure the rationality of its collective judgments. But this conclusion would be too quick. The impossibility theorem should be seen as characterizing the logical space of aggregation procedures (LIST AND PETTIT [2002]; LIST [2006]). In particular, we can characterize different aggregation procedures in terms of which conditions they meet and which they violate.

If a group acting as a distributed cognitive system seeks to ensure the rationality of its collective judgments, the group must use an aggregation procedure that violates at least one of the conditions of the theorem.

4.3 First solution: relaxing universal domain

If the amount of disagreement in a particular group is limited or if the group has mechanisms in place for reducing disagreement – such as mechanisms of group deliberation – the group might use an aggregation procedure that violates universal domain. For example, a deliberating group that successfully avoids combinations of individual judgments of the kind in table 2 might use majority voting as its aggregation procedure and yet generate rational collective judgments.

But this solution does not work in general. Even in an expert committee whose task is to make judgments on factual matters without conflicts of interest, disagreement may still be significant and pervasive. Although one can study conditions that make the occurrence of judgment combinations of the kind in table 2 less likely (DRYZEK AND LIST [2003]; LIST [2002]), I set this issue aside here and assume that groups that are faced with primarily cognitive tasks (as opposed to primarily political ones, for example) should normally use aggregation procedures satisfying universal domain.

4.4 Second solution: relaxing anonymity

It can be shown that, if anonymity is relaxed but the other two conditions are retained, the only possible aggregation procedure is a 'dictatorial procedure', whereby the collective judgments are always those of some antecedently fixed group member (the 'dictator') (PAULY AND VAN HEES [2005]). Some groups might put one individual – say a committee chair – in charge of forming its collective judgments. But this solution clearly conflicts with the idea of collectively distributed cognition, and as discussed below, a group organized in this dictatorial way loses out on the epistemic advantages of distributed cognition.

However, below I also suggest that a group acting as a distributed cognitive system may sometimes benefit from relaxing anonymity together with systematicity and implementing a division of cognitive labour whereby different components of a complex cognitive task are allocated to different subgroups.

4.5 Third solution: relaxing systematicity

A potentially promising solution lies in relaxing systematicity, i.e. treating different propositions differently in the process of forming collective judgments. For the purposes of a given cognitive task, a group may designate some propositions as 'premises' and others as 'conclusions' and assign epistemic priority either to the premises or to the conclusions (for a more extensive discussion of this process, see LIST [2006]).

If the group assigns priority to the premises, it may use the so-called 'premise-based procedure', whereby the group first makes a collective judgment on each premise by taking a majority vote on that premise and then derives its collective judgments on the conclusions from these collective judgments on the premises. In the expert committee example, propositions p and $p{\rightarrow}q$ might be designated as premises (perhaps on the grounds that p and $p{\rightarrow}q$ are more basic than q), and proposition q might be designated as a conclusion. The committee might then take majority votes on p and $p{\rightarrow}q$ and derive its judgment on q from its judgments on p and $p{\rightarrow}q$.[8]

Alternatively, if the group assigns priority to the conclusions, it may use the so-called 'conclusion-based procedure', whereby the group takes a majority vote only on each conclusion and makes no collective judgments on the premises. In addition to violating systematicity, this aggregation procedure fails to produce complete collective judgments. But sometimes a group is required to make judgments only on conclusions, but not on premises, and in such cases incompleteness in the collective judgments on the premises may be defensible.

The premise- and conclusion-based procedures are not the only aggregation procedures violating systematicity. Further important possibilities arise when both systematicity and anonymity are relaxed. The group can then use an aggregation procedure that not only assigns priority to the premises, but also assigns different such premises to different subgroups and thereby implements a particularly clear form of distributed cognition. Specifically, the group may use the so-called 'distributed premise-based procedure'. Here different individuals specialize on different premises and give their individual judgments only on these premises. Now the group makes a

---

[8] In the present example, the truth-value of q is not always settled by the truth-values of p and $p{\rightarrow}q$; so the group may need to stengthen its premises in order to make them sufficient to determine its judgment on the conclusion.

collective judgment on each premise by taking a majority vote on that premise among the relevant 'specialists', and then the group derives its collective judgments on the conclusions from these collective judgments on the premises. This procedure is discussed in greater detail below.

For many cognitive tasks performed by groups, giving up systematicity and using a (regular or distributed) premise-based or conclusion-based procedure may be an attractive way to avoid the impossibility result explained above. Each of these procedures allows a group to produce rational collective judgments. Arguably, a premise-based or distributed premise-based procedure makes the group's performance as a unified cognitive system particularly visible. A group using such a procedure acts as a reason-driven system when it derives its collective judgments on conclusions from its collective judgments on relevant premises.

However, giving up systematicity comes with a price. Aggregation procedures that violate systematicity may be vulnerable to manipulation by prioritizing propositions strategically, and strategic agents with agenda-setting influence over the group might exploit these strategic vulnerabilities.

For example, in the case of a regular premise-based procedure, the collective judgments may be sensitive to the choice of premises. In the example of table 2, if p and $p{\rightarrow}q$ are designated as premises, then all three propositions, p, $p{\rightarrow}q$ and q, are collectively judged to be true. If p and q are designated as premises, then p is judged to be true and both q and $p{\rightarrow}q$ are judged to be false; finally, if q and $p{\rightarrow}q$ are designated as premises, then $p{\rightarrow}q$ is judged to be true, and both p and q are judged to be false. Although there seems to be a natural choice of premises in the present example, namely p and $p{\rightarrow}q$, this may not generally be the case, and the outcome of a premise-based procedure may therefore depend as much on the choice of premises as it depends on the individual judgments to be aggregated. In the present example, an environmental activist may prefer to prioritize the propositions in such a way as to bring about the collective judgment that proposition q is true, while a transport lobbyist may prefer to prioritize them in such a way as to bring about the opposite judgment on q.

Under the distributed premise-based procedure, an additional sensitivity to the choice of 'specialists' on each premise arises. Likewise, in the case of the conclusion-based procedure, the choice of conclusions obviously matters, since the group makes collective judgments only on these conclusions and on no other propositions.[9]

4.6 Fourth solution: permitting incomplete collective judgments

The first three solutions proposed in response to the impossibility theorem above have required relaxing one of the three minimal conditions on how individual judgments are aggregated into collective judgments. The present solution preserves these minimal conditions, but weakens the requirements on the collective judgments themselves by permitting incompleteness in these judgments (see also LIST [2006]).

If a group acting as an overall cognitive system is prepared to refrain from making a collective judgment on some propositions – namely on those on which there is too much disagreement between the group members – then it may use an aggregation

---

[9] It can be shown that in some important respects, the premise-based procedure is more vulnerable to strategic manipulation than the conclusion-based procedure. See DIETRICH AND LIST [2005b].

procedure such as the 'unanimity procedure', whereby the group makes a judgment on a proposition if and only if the group members unanimously endorse that judgment. Propositions judged to be true by all members are collectively judged to be true; and ones judged to be false by all members are collectively judged to be false; no collective judgment is made on any other propositions. (Instead of the unanimity procedure, the group might also use 'supermajority voting' with a sufficiently large supermajority threshold.)

Groups operating in a strongly consensual manner may well opt for this solution, but in many cases making no judgment on some propositions is simply not an option. For example, when an expert committee is asked to give advice on a particular issue, it is usually expected to take a determinate stance on that issue.

4.7 Lessons to be drawn

I have shown that a group's capacity to form rational collective judgments depends on the group's aggregation procedure: a group acting as a distributed cognitive system can ensure the rationality of its collective judgments on some non-trivially interconnected propositions only if it uses a procedure that violates one of universal domain, anonymity or systematicity or that produces incomplete collective judgments. Moreover, different aggregation procedures may lead to different collective judgments for the same combination of individual judgments. As an illustration, table 3 shows the collective judgments for the individual judgments in table 2 under different aggregation procedures.

Table 3: Different aggregation procedures applied to the individual judgments in table 2

	p	$p{\to}q$	q
Majority voting*	True	True	False
Premise-based procedure with p, $p{\to}q$ as premises	True	True	True
Conclusion-based procedure with q as conclusion	No judgment	No judgment	False
Distributed premise-based procedure with individual 1 specializing on p and individual 2 specializing on $p{\to}q$	True	False	False
Unanimity procedure	No judgment	No judgment	No judgment
Dictatorship of individual 3	False	True	False

* inconsistent

If we were to assess a group's performance as a distributed cognitive system solely on the basis of whether the group's collective judgments are rational, this would give us insufficient grounds for selecting a unique aggregation procedure. As I have illustrated, many different aggregation procedures generate consistent collective judgments, and even if we require completeness in addition to consistency, several possible aggregation procedures remain. To recommend a suitable aggregation procedure that a group can use for a given cognitive task, the question of whether the group produces rational collective judgments is, by itself, not a sufficient criterion.

5 Truth-tracking in a distributed cognitive system

Can a group acting as a distributed cognitive system generate collective judgments that track the truth? Following NOZICK [1981], a system 'tracks the truth' on some proposition p if two conditions are met. First, if – actually or counterfactually – p were true, the system would judge p to be true. Second, if – actually or counterfactually – p were not true, the system would not judge p to be true. These conditions can be applied to any cognitive system, whether it consists just of a single agent or of multiple agents acting together. In particular, if a group's organizational structure allows the group to form collective judgments, then one can ask whether these judgments satisfy Nozick's two conditions.

As a simple measure of how well a system satisfies Nozick's two conditions, I consider two conditional probabilities (LIST [2006]): the probability that the system judges p to be true given that p is true, and the probability that the system does not judge p to be true given that p is false. Call these two conditional probabilities the system's 'positive' and 'negative reliability' on p, respectively.

By considering a group's positive and negative reliability on various propositions under different aggregation procedures and different scenarios, I now show that it is possible for a group acting as a distributed cognitive system to track the truth, but that, once again, the aggregation procedure affects the group's success.

5.1 The first scenario and its lesson: epistemic gains from democratization

Suppose that a group is given the cognitive task of making a collective judgment on a single factual proposition, such as proposition p in the expert committee example above. As a baseline scenario (e.g. GROFMAN, OWEN AND FELD [1983]), suppose that the group members hold individual judgments on proposition p, where two conditions are met. First, each group member has the same positive and negative reliability r on proposition p, where $1 > r > 1/2$ (the 'competence' condition); so individual judgments are noisy but biased towards the truth. Second, the judgments of different group members are mutually independent (the 'independence' condition). (Obviously, it is also important to study scenarios where these conditions are violated, and below I consider some such scenarios.[10])

A group acting as a distributed cognitive system must use an aggregation procedure to make its collective judgment on p based on the group members' individual judgments on p. What is the group's positive and negative reliability on p under different aggregation procedures?

Let me compare three different procedures: first, a dictatorial procedure, where the collective judgment is always determined by the same fixed group member; second, the unanimity procedure, where agreement among all group members is necessary for reaching a collective judgment; and third, majority voting, which perhaps best implements the idea of a democratically organized form of distributed cognition (at least in the case of a single proposition).

---

[10] Cases where different individuals have different levels of reliability are discussed, for example, in GROFMAN, OWEN AND FELD [1983] and BORLAND [1989]. Cases where there are dependencies between different individuals' judgments are discussed, for example, in LADHA [1992], ESTLUND [1994] and DIETRICH AND LIST [2004]. Cases where individuals express their judgments strategically rather than truthfully are discussed in AUSTEN-SMITH AND BANKS [1996].

Under a dictatorial procedure, the group's positive and negative reliability on p equals that of the dictator, which is r by assumption.

Under the unanimity procedure, the group's positive reliability on p equals r^n, which approaches 0 as the group size increases, but its negative reliability on p equals $1-(1-r)^n$, which approaches 1 as the group size increases. This means that the unanimity procedure is good at avoiding false positive judgments, but bad at reaching true positive ones. A determinate collective judgment on p is reached only if all individuals agree on the truth-value of p; if they don't agree, no collective judgment on p is made.

Finally, under majority voting, the group's positive and negative reliability on p approaches 1 as the group size increases. Why does this result hold? Each individual has a probability $r>0.5$ of making a correct judgment on p; by the law of large numbers, the proportion of individuals who make a correct judgment on p approaches $r>0.5$ as the group size increases and thus constitutes a majority with a probability approaching 1. Informally, majority voting allows the group to extract the signal from the group members' judgments, while filtering out the noise. This is the famous 'Condorcet jury theorem' (e.g. GROFMAN, OWEN AND FELD [1983]).

Table 4 shows the group's positive and negative reliability on p under majority voting and under a dictatorial procedure, and tables 5 and 6 show, respectively, the group's positive and negative reliability on p under a dictatorial procedure and under the unanimity procedure. In each case, individual group members are assumed (as an illustration) to have a positive and negative reliability of $r=0.54$ on p. In all tables, the group size is on the horizontal axis and the group's reliability on the vertical axis.[11]

Table 4: The group's positive and negative reliability on p: majority voting (top curve); dictatorship (bottom curve) (setting r=0.54 as an illustration)

---
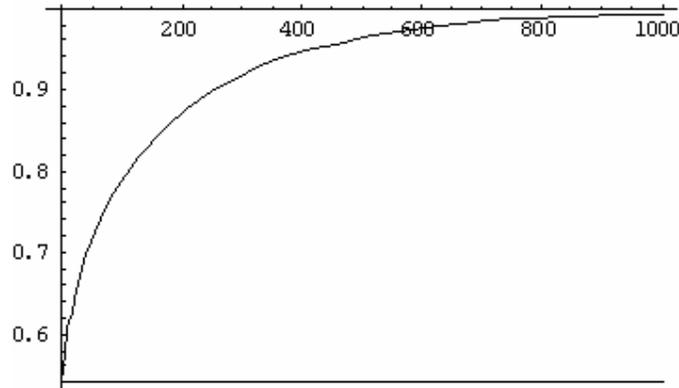
[11] The present curves are the result of averaging between two separate curves for even- and odd-numbered group sizes. When the group size is an even number, the group's reliability may be lower because of the possibility of majority ties.

Table 5: The group's positive reliability on p: dictatorship (top curve); unanimity procedure (bottom curve) (setting r = 0.54 as an illustration)
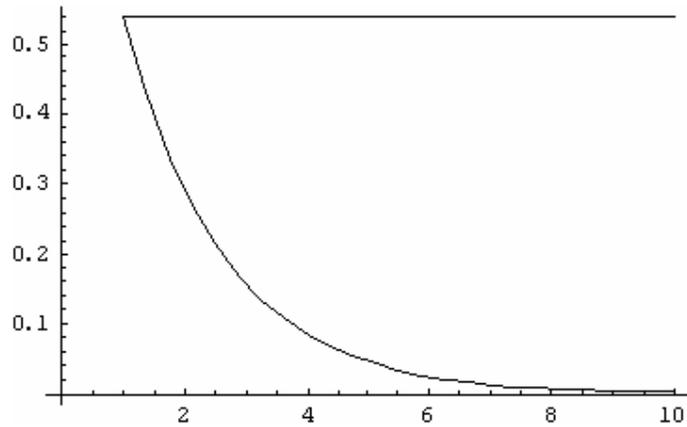
Table 6: The group's negative reliability on p: unanimity procedure (top curve); dictatorship (bottom curve) (setting r = 0.54 as an illustration)
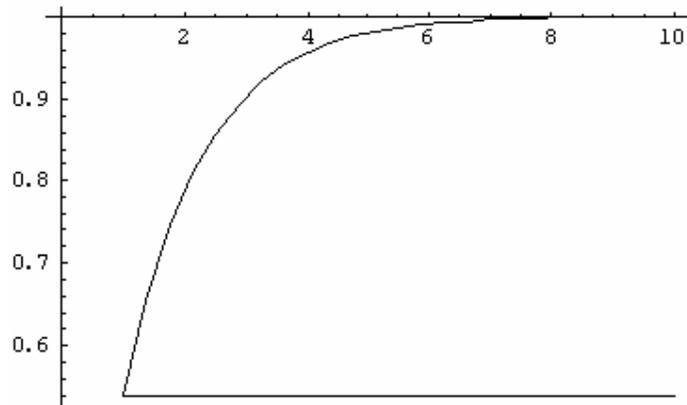
What lessons can be drawn from this first scenario? If individuals are independent, fallible, but biased towards the truth, majority voting outperforms both dictatorial and unanimity procedures in terms of maximizing the group's positive and negative reliability on *p*. The unanimity procedure is attractive only in those special cases where the group seeks to minimize the risk of making false positive judgments, such as in some jury decisions. A dictatorial procedure fails to pool the information held by different individuals.

Hence, when a group acting as a distributed cognitive system seeks to track the truth, there may be 'epistemic gains from democratization', i.e. from making a collective judgment on a given proposition democratically by using majority voting. More generally, even when individual reliability differs between individuals, a weighted form of majority voting still outperforms a dictatorship by the most reliable individual: each individual's vote simply needs to have a weight proportional to $log(r/(1-r))$, where *r* is the individual's reliability on the proposition in question (BEN-YASHAR AND NITZAN [1997]).

5.2 The second scenario and its lesson: epistemic gains from disaggregation

Suppose now that a group is given the cognitive task of making a collective judgment not only on a single factual proposition, but on a set of interconnected factual propositions. As an illustration, suppose that there are $k>1$ premises p_1, ..., p_k and a conclusion q, where q is true if and only if the conjunction of p_1, ..., p_k is true. This structure also allows representing a variant of the expert committee example above. For extensive discussions of the present scenario and other related scenarios, see BOVENS AND RABINOWICZ [2005] and LIST [2005a], [2006]. Analogous points apply to the case where q is true if and only if the disjunction of p_1, ..., p_k is true.

In this case of multiple interconnected propositions, individuals cannot generally have the same reliability on all propositions. Suppose, as an illustration, that each individual has the same positive and negative reliability r on each premise p_1, ..., p_k and makes independent judgments on different premises. Then each individual's positive reliability on the conclusion q is r^k, which is below r and often below 0.5 (whenever $r < k$-th root of 0.5), while his or her negative reliability on q is above r. Here individuals are much worse at detecting the truth of the conclusion than the truth of each premise, but much better at detecting the falsehood of the conclusion than the falsehood of each premise. In the expert committee example, it might be easier to make correct judgments on propositions p and $p{\rightarrow}q$ than on proposition q. Of course, other scenarios can also be constructed, but the point remains that individuals typically have different levels of reliability on different propositions (LIST [2006]).

What is the group's positive and negative reliability on the various propositions under different aggregation procedures? As before, suppose the judgments of different group members are mutually independent.

Majority voting performs well only on those propositions on which individuals have a positive and negative reliability above 0.5. As just argued, individuals may not meet this condition on all propositions. Moreover, majority voting does not generally produce consistent collective judgments (on the probability of majority inconsistencies, see LIST [2005a]). Let me now compare dictatorial, conclusion-based and premise-based procedures.

Under a dictatorial procedure, the group's positive and negative reliability on each proposition equals that of the dictator; in particular, the probability that *all* propositions are judged correctly is r^k, which may be very low, especially when the number of premises k is large.

Under the conclusion-based procedure, unless individuals have a high reliability on each premise, namely $r > k$-th root of 0.5 (e.g. 0.71 when $k=2$, or 0.79 when $k=3$), the group's positive reliability on the conclusion q approaches 0 as the group size increases. Its negative reliability on q approaches 1. Like the unanimity procedure in the single-proposition case, the conclusion-based procedure is good at avoiding false positive judgments on the conclusion, but (typically) bad at reaching true positive ones (see also BOVENS AND RABINOWICZ [2005]).

Under the premise-based procedure, the group's positive and negative reliability on every proposition approaches 1 as the group size increases. This result holds because, by the Condorcet jury theorem as stated above, the group's positive and negative reliability on each premise p_1, ..., p_k approaches 1 with increasing group size, and therefore the probability that the group derives a correct judgment on the conclusion also approaches 1 with increasing group size.

As illustration, suppose that there are $k=2$ premises and individuals have a positive and negative reliability of $r=0.54$ on each premise. Table 7 shows the group's probability of judging *all* propositions correctly under the premise-based procedure and under a dictatorial procedure. Tables 8 and 9 show, respectively, the group's positive and negative reliability on the conclusion q under a dictatorial procedure and under the conclusion-based procedure.

Table 7: The group's probability of judging all propositions correctly: premise-based procedure (top curve); dictatorship (bottom curve) (setting r = 0.54 as an illustration)
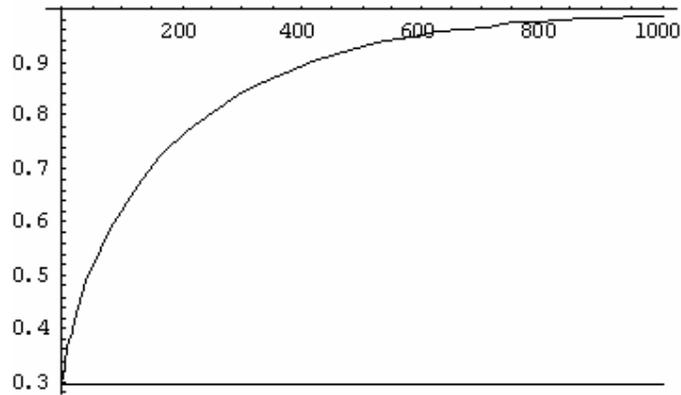

Table 8: The group's positive reliability on the conclusion q: dictatorship (top curve); conclusion-based procedure (bottom curve) (setting r = 0.54 as an illustration)
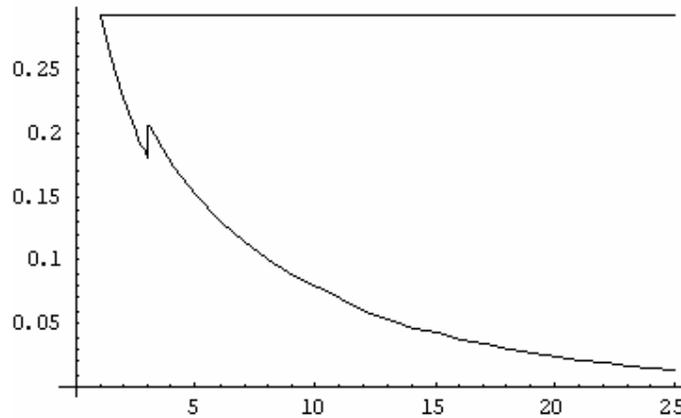
Table 9: The group's negative reliability on the conclusion q: conclusion-based procedure (top curve); dictatorship (bottom curve) (setting r = 0.54 as an illustration)
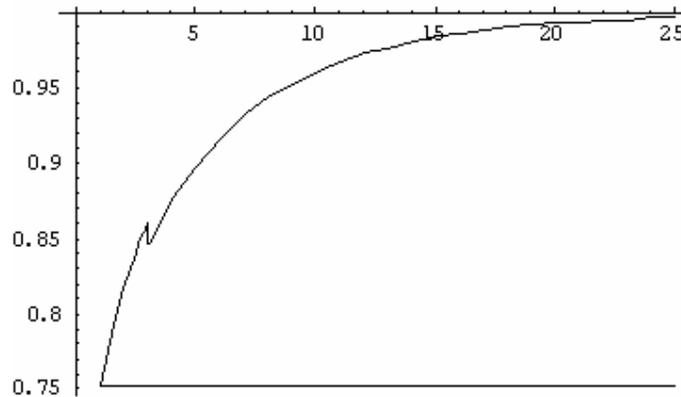
What lessons can be drawn from this second scenario? Under the present assumptions, the premise-based procedure outperforms both dictatorial and conclusion-based procedures in terms of simultaneously maximizing the group's positive and negative reliability on every proposition. Like the unanimity procedure before, the conclusion-based procedure is attractive only when the group seeks to minimize the risk of making false positive judgments on the conclusion; again, a dictatorial procedure is bad at information pooling.

Hence, if a larger cognitive task such as making a judgment on some conclusion can be disaggregated into several smaller cognitive tasks such as making judgments on relevant premises, then there may be 'epistemic gains from disaggregation', i.e. from making collective judgments on that conclusion on the basis of separate collective judgments on those premises. (For further results and a discussion of different scenarios, see BOVENS AND RABINOWICZ [2005] and LIST [2006].)

5.3 The third scenario and its lesson: epistemic gains from distribution

When a group is faced with a complex cognitive task that requires making judgments on several propositions, different members of the group may have different levels of expertise on different propositions. This is an important characteristic of many committees, groups of scientific collaborators, large organizations, and so on. Moreover, each individual may lack the temporal, computational and informational resources to become sufficiently reliable on every proposition. If we take this problem into account, can we improve on the premise-based procedure?

Suppose, as before, that a group has to make collective judgments on $k>1$ premises p_1, ..., p_k and a conclusion q, where q is true if and only if the conjunction of p_1, ..., p_k is true. Instead of requiring every group member to make a judgment on every premise, we might partition the group into k subgroups (for simplicity, of approximately equal size), where the members of each subgroup specialize on one premise and make a judgment on that premise alone. Instead of a using a regular premise-based procedure as in the previous scenario, the group might now use a distributed premise-based procedure: the collective judgment on each premise is made by taking a majority vote within the subgroup specializing on that premise, and the collective judgment on the conclusion is then derived from these collective judgments on the premises.

When does the distributed premise-based procedure outperform the regular premise-based procedure at maximizing the group's probability of making correct judgments on the propositions?

Intuitively, there are two effects here that pull in opposite directions. First, there may be 'epistemic gains from specialization': individuals may become more reliable on the proposition on which they specialize. But, second, there may also be 'epistemic losses from lower numbers': each subgroup voting on a particular proposition is smaller than the original group (it is only approximately $1/k$ the size of the original group when there are k premises), which may reduce the benefits from majoritarian judgment aggregation on that proposition.

Whether or not the distributed premise-based procedure outperforms the regular premise-based procedure depends on which of these two opposite effects is stronger. Obviously, if there were no epistemic gains from specialization, then the distributed premise-based procedure would suffer only from losses from lower numbers on each premise and would therefore perform worse than the regular premise-based procedure. On the other hand, if the epistemic losses from lower numbers were relatively small compared to the epistemic gains from specialization, then the distributed premise-based procedure would outperform the regular one. The following result holds:

Theorem. For any group size n (divisible by k), there exists an individual (positive and negative) reliability level $r^* > r$ such that the following holds: if, by specializing on some proposition p, individuals achieve a reliability above r^* on p, then the majority judgment on p in a subgroup of n/k specialists (each with reliability r^* on p) is more reliable than the majority judgment on p in the original group of n non-specialists (each with reliability r on p).

Hence, if by specializing on one premise, individuals achieve a reliability above r^* on that premise, then the distributed premise-based procedure outperforms the regular premise-based procedure. How great must the reliability increase from r to r^* be to have this effect? Strikingly, a small reliability increase typically suffices. Table 10 shows some sample calculations. For example, when there are $k=2$ premises, if the original individual reliability was $r=0.52$, then a reliability above $r^*=0.5281$ after specialization suffices; it it was $r=0.6$, then a reliability above $r^*=0.6393$ after specialization suffices.

Table 10: Reliability increase from r to r* required to outweigh the loss from lower numbers

	$k = 2, n = 50$			$k = 3, n = 51$			$k = 4, n = 52$		
$r =$	0.52	0.6	0.75	0.52	0.6	0.75	0.52	0.6	0.75
$r^* =$	0.5281	0.6393	0.8315	0.5343	0.6682	0.8776	0.5394	0.6915	0.9098

Table 11 shows the group's probability of judging *all* propositions correctly under regular and distributed premise-based procedures, where there are $k=2$ premises and where individuals have positive and negative reliabilities of $r=0.54$ and $r^*=0.58$ before and after specialization, respectively.

Table 11: The group's probability of judging all propositions correctly: distributed (top curve) and regular premise-based procedure (bottom curve) (setting r = 0.54 and r* = 0.58 as an illustration)
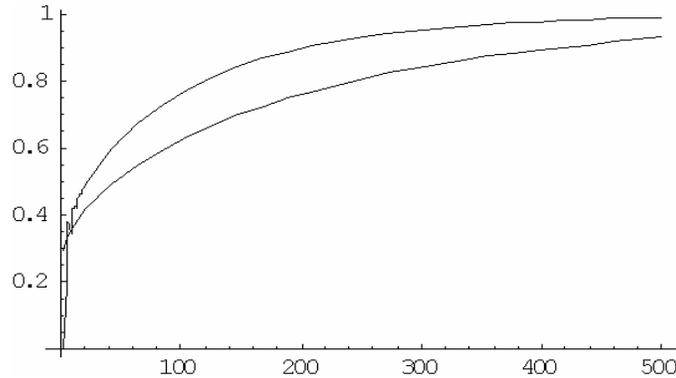
What lessons can be drawn from this third scenario? Even when there are only relatively modest gains from specialization, the distributed premise-based procedure may outperform the regular premise-based procedure in terms of maximizing the group's positive and negative reliability on every proposition.

Hence there may be 'epistemic gains from distribution': if a group has to perform a complex cognitive task, the group may benefit from subdividing the task into several smaller tasks and distributing these smaller tasks across multiple subgroups. Plausibly, such division of cognitive labour is the mechanism underlying the successes of collectively distributed cognition in science, as investigated by KNORR CETINA [1999], GIERE [2002] and others. The research practices in large-scale collaborative research projects, such as those in high-energy physics or in other large expert teams as mentioned above, rely on mechanisms similar to those represented, in a stylized form, by the distributed premise-based procedure.

In conclusion, a group acting as a distributed cognitive system can succeed at tracking the truth, but the group's aggregation procedure plays an important role in determining the group's success.

6 Concluding remarks

I have discussed collectively distributed cognition from a social-choice-theoretic perspective. In particular, I have introduced the emerging theory of judgment aggregation to propose a way of modelling a group as a distributed cognitive system, i.e. as a system that can generate collective judgments. Within this framework, I have asked whether such a group can be rational and track the truth in its collective judgments. My main finding is that a group's performance as a distributed cognitive system depends crucially on its aggregation procedure, and I have investigated *how* the aggregation procedure matters.

With regard to a group's rationality as a distributed cognitive system, I have discussed an impossibility theorem by which we can characterize the logical space of aggregation procedures that a group can use to generate rational collective judgments. No aggregation procedure generating consistent and complete collective judgments can simultaneously satisfy universal domain, anonymity and systematicity. To find an aggregation procedure that produces rational collective judgments, it is therefore necessary to relax one of universal domain, anonymity or systematicity, or to weaken

the requirement of rationality itself by permitting incomplete collective judgments. Which relaxation is most defensible depends on the group and cognitive task in question.

With regard to a group's capacity to track the truth as a distributed cognitive system, I have identified three effects that are relevant to the design of a good aggregation procedure: there may be epistemic gains from democratization, disaggregation and distribution. Again, the applicability and magnitude of each effect depends on the group and cognitive task in question, and there may not be a 'one size fits all' aggregation procedure that is best for all groups and all cognitive tasks. But the fact that a group may sometimes benefit from the identified effects reinforces the potential of epistemic gains through collectively distributed cognition.

The present results give a fairly optimistic picture of a group's capacity to perform as a distributed cognitive system. I have thereby focused on cooperative rather than competitive practices within groups or communities. It is an important empirical question how pervasive such cooperative practices are and how often the favourable conditions such practices require are met. Clearly, scientific communities are characterized by both competitive and cooperative practices. Much research in the sociology and economics of science has focused on competitive practices (as evidenced by the theme of the 2005 Conference on New Political Economy). There has also been much research on rationality failures and inefficiencies that can arise in groups trying to perform certain tasks at a collective level. Public choice theorists, in particular, have highlighted the impossibility results on democratic aggregation and the pervasiveness of suboptimal equilibria in various collective interactions.

Clearly, the details of my rather more optimistic results depend on various assumptions and may change with changes in these assumptions. But my aim has not been to argue that *all* groups acting as distributed cognitive systems perform well; indeed, this claim is likely to be false. Rather, my aim has been to show that successful distributed cognition in a group is *possible* and to illustrate the usefulness of the theory of judgment aggregation for investigating how it is possible and under what conditions.

References

ARROW, KENNETH J. [1951], *Social Choice and Individual Values*, Wiley: New York.

AUSTEN-SMITH, DAVID, AND JEFFREY BANKS [1996], "Information Aggregation, Rationality and the Condorcet Jury Theorem," *American Political Science Review*, 90, 34-45.

BEN-YASHAR, RUTH, AND SHMUEL NITZAN [1997], "The optimal decision rule for fixed-size committees in dichotomous choice situations: the general result," *International Economic Review*, 38, 175-186.

BORLAND, PHILIP [1989], "Majority systems and the Condorcet jury theorem," *Statistician*, 38, 181-189.

BOVENS, LUC, AND WLODEK RABINOWICZ [2005], "Democratic Answers to Complex Questions: an Epistemic Perspective," *Synthese*, forthcoming.

DENNETT, DANIEL [1987], *The Intentional Stance*, MIT Press: Cambridge, MA.

DIETRICH, FRANZ [2005], "Judgment aggregation: (im)possibility theorems," *Journal of Economic Theory*, forthcoming.

DIETRICH, FRANZ, AND CHRISTIAN LIST [2004], "A model of jury decisions where all jurors have the same evidence," *Synthese*, 142, 175-202.

DIETRICH, FRANZ, AND CHRISTIAN LIST [2005a], "Arrow's theorem in judgment aggregation," Working paper, London School of Economics, London.

DIETRICH, FRANZ, AND CHRISTIAN LIST [2005b], "Strategy-proof judgment aggregation," Working paper, London School of Economics, London.

DRYZEK, JOHN, AND CHRISTIAN LIST [2003], "Social Choice Theory and Deliberative Democracy: A Reconciliation," *British Journal of Political Science*, 33(1), 1-28.

ESTLUND, DAVID [1994], "Opinion leaders, independence and Condorcet's jury theorem," *Theory and Decision*, 36, 131-162.

GIERE, RONALD [2002], "Distributed Cognition in Epistemic Cultures," *Philosophy of Science*, 69, 637-644.

GOLDMAN, ALVIN [1999], *Knowledge in a Social World*, Oxford University Press: Oxford.

GOLDMAN, ALVIN [2004], "Group Knowledge versus Group Rationality: Two Approaches to Social Epistemology," *Episteme: A Journal of Social Epistemology*, 1(1), 11-22.

GOODIN, ROBERT E. (ed.) [1996], *The Theory of Institutional Design*, Cambridge University Press: Cambridge.

GROFMAN, BERNARD, GUILLERMO OWEN AND SCOTT FELD [1983], "Thirteen theorems in search of the truth," *Theory and Decision*, 15, 261-278.

HUTCHINS, EDWIN [1995], *Cognition in the Wild*, MIT Press: Cambridge, MA.

KNORR CETINA, KARIN [1999], *Epistemic Cultures: How the Sciences Make Knowledge*, Harvard University Press: Cambridge, MA.

KORNHAUSER, LEWIS A., AND LAWRENCE G. SAGER [1986], "Unpacking the Court," *Yale Law Journal*, 96(1), 82-117.

LADHA, KRISHNA [1992], "The Condorcet jury theorem, free speech and correlated votes," *American Journal of Political Science*, 36, 617-634.

LIST, CHRISTIAN [2002], "Two Concepts of Agreement," *The Good Society*, 11(1), 72-79.

LIST, CHRISTIAN [2005a], "The Probability of Inconsistencies in Complex Collective Decisions," *Social Choice and Welfare*, 24(1), 3-32.

LIST, CHRISTIAN [2005b], "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective," *Episteme: A Journal of Social Epistemology*, 2(1), 25-38.

LIST, CHRISTIAN [2006], "The Discursive Dilemma and Public Reason," *Ethics*, forthcoming.

LIST, CHRISTIAN, AND PHILIP PETTIT [2002], "Aggregating Sets of Judgments: An Impossibility Result," *Economics and Philosophy*, 18, 89-110.

LIST, CHRISTIAN, AND PHILIP PETTIT [2004], "Aggregating Sets of Judgments: Two Impossibility Results Compared," *Synthese*, 140(1-2), 207-235.

LIST, CHRISTIAN, AND PHILIP PETTIT [2005a], "On the Many as One," *Philosophy and Public Affairs*, 33(4), 377-390.

LIST, CHRISTIAN, AND PHILIP PETTIT [2005b], "Group Agency and Supervenience," *Southern Journal of Philosophy*, forthcoming.

NOZICK, ROBERT [1981], *Philosophical Explanations*, Harvard University Press: Cambridge, MA.

PAULY, MARC, AND MARTIN VAN HEES [2005], "Logical Constraints on Judgment Aggregation," *Journal of Philosophical Logic*, forthcoming.

PETTIT, PHILIP [2001], "Deliberative Democracy and the Discursive Dilemma," *Philosophical Issues* (supplement to *Nous*), 11, 268-99.

PETTIT, PHILIP [2003], "Groups with Minds of their Own," in FRED SCHMITT (ed.), *Socializing Metaphysics*, Rowan and Littlefield: New York.

ROVANE, CAROL [1998], *The Bounds of Agency: An essay in Revisionary Metaphysics*, Princeton University Press: Princeton.

Christian List
Department of Government
London School of Economics
London WC2A 2AE
United Kingdom
c.list@lse.ac.uk