

The Discursive Dilemma and Public Reason

Christian List¹
forthcoming in *Ethics*

Political theorists have offered many accounts of collective decision-making under pluralism. I discuss a key dimension on which such accounts differ: the importance assigned not only to the choices made but also to the reasons underlying those choices. On that dimension, different accounts lie in between two extremes. The ‘minimal liberal account’ holds that collective decisions should be made only on practical actions or policies and that underlying reasons should be kept private. The ‘comprehensive deliberative account’ stresses the importance of giving reasons for collective decisions, where such reasons should also be collectively decided. I compare these two accounts on the basis of a formal model developed in the growing literature on the ‘discursive dilemma’ and ‘judgment aggregation’ and address several questions: What is the trade-off between the (minimal liberal) demand for reaching agreement on outcomes and the (comprehensive deliberative) demand for reason-giving? How large should the ‘sphere of public reason’ be? When do the decision procedures suggested by the two accounts agree and when not? How good are these procedures at truth-tracking on factual matters? What strategic incentives do they generate for decision-makers? My discussion identifies what is at stake in the choice between minimal liberal and comprehensive deliberative accounts of collective decision-making, and sheds light not only on these two ideal-typical accounts themselves, but also on many characteristics that intermediate accounts share with them.

I. INTRODUCTION

Collective decisions are ubiquitous in complex democratic societies. Elections, referenda, decisions in legislatures, committees, multi-member courts, expert panels and boards of companies or other organizations are all examples. In such decisions, disagreement is equally ubiquitous. People disagree with each other on many levels. They disagree not only on *what* choices should be made, but also on *why* those choices should be made.

Political theorists have offered different accounts of how collective decisions should be made under conditions of pluralism and when such decisions are legitimate. Obviously, different decision problems may require different decision procedures or different criteria of legitimacy. But, even for a given decision problem, rival accounts of how to solve the problem are usually on offer, which differ on several dimensions. In this article, I focus on one such dimension: the importance assigned not only to the *choices* made (the ‘what’ question) but also to the *reasons* underlying those choices (the ‘why’ question). This dimension can be seen as a spectrum between two extremes: the *minimal liberal* account at one end, and the *comprehensive deliberative* account at the other.

The minimal liberal account emphasizes the ‘what’ question, the comprehensive deliberative account the ‘why’ question. The minimal liberal account holds that collective

¹ Department of Government, London School of Economics and Political Science, London WC2A 2AE, U.K. An earlier version of this article was presented at the Colloquium on Law, Economics and Politics, NYU Law School, September 2004. I am grateful to the participants at this Colloquium for a challenging and helpful discussion. I especially thank Steven Brams, Franz Dietrich, John Ferejohn, Lewis Kornhauser, Philip Pettit and the editors and referees of *Ethics* for helpful comments, suggestions and editorial advice.

decisions should be made only on practical actions or policies and that the reasons (or justifications) underlying those decisions should be kept private. The comprehensive deliberative account, by contrast, stresses the importance of giving reasons for collective decisions, where those reasons should themselves be collectively decided.

Consider a decision to protect a natural habitat. On the minimal liberal account, it is sufficient to reach agreement on a protection policy, without agreeing on the reasons supporting that policy. On the comprehensive deliberative account, these reasons should also be publicly considered and decided. There could be many such reasons, and despite agreeing that the habitat should be protected, people may still disagree on whether this should be done for aesthetic, recreational, economic, religious or other reasons.

While these two extreme positions are ideal types, several influential accounts, including variants of Dworkin's, Pettit's, Rawls's and Sunstein's accounts, can be placed at more intermediate positions between them, as discussed below. One might also defend a minimal liberal account (or something similar) for some decision problems while defending a comprehensive deliberative account (or something similar) for others.

My aim in this article is to compare the minimal liberal and comprehensive deliberative accounts on the basis of a formal model developed in the growing literature on the 'discursive dilemma' and 'judgment aggregation'.² This model can represent the two accounts in a unified framework, albeit in a simplified form. I address several questions: What is the trade-off between the (minimal liberal) demand for reaching agreement on outcomes and the (comprehensive deliberative) demand for reason-giving? How large should the 'sphere of public reason' be? When do the decision procedures suggested by the two accounts agree and when do they conflict? How good are these procedures at tracking the truth in decisions involving matters of fact? What strategic incentives do they generate for decision-makers? Answers to these questions allow us to

² The 'discursive dilemma' was first discussed in a legal context under the name 'doctrinal paradox' by Kornhauser and Sager (1986, 1993, 2004), Kornhauser (1992) and Chapman (1998). Brennan (2001) and Pettit (2001a), who introduced the name 'discursive dilemma', reinterpreted the 'doctrinal paradox' as a broader problem of 'collective (in)coherence'. Generalizing the 'discursive dilemma' by drawing on Arrow's (1951) axiomatic approach to social choice theory, List and Pettit (2002, 2004) have developed a formal model of judgment aggregation and proved a simple impossibility theorem. Further technical results and generalizations of the model have been obtained by Pauly and van Hees (2005), Dietrich (2005a, 2005b), List (2003, 2004), Dietrich and List (2004a, 2005), Pigozzi (2004), Gärdenfors (2005), Nehring and Puppe (2005) and Dokow and Holzman (2005). Bovens and Rabinowicz (2005) and List (2005) have discussed the 'discursive dilemma' from the perspective of the Condorcet jury theorem.

compare the strengths and weaknesses of the two accounts in relation to different decision problems.

After introducing the two accounts in Section II, I argue in Section III that they can be associated with different decision procedures – ‘conclusion-’ and ‘premise-based’ ones – which may generate different outcomes for the same decision problem: the ‘discursive dilemma’. In Section IV, I introduce the concept of the public agenda, and in Section V, I identify a trade-off between some general demands on collective decision-making and assess how the conclusion- and premise-based procedures solve the trade-off. Sections VI and VII address, respectively, the truth-tracking capacities of the two procedures in factual decisions and the incentives they generate for decision makers. In Section VIII, I conclude and ask whether the premise- and conclusion-based procedures can be embedded in a unified account of collective decision-making that constitutes a compromise between the minimal liberal and comprehensive deliberative extremes.

Although the two ideal-typical accounts discussed here are less subtle than many accounts defended in the literature, a comparison of the ideal types can still shed light on many characteristics that intermediate accounts have in common with them.

II. TWO ACCOUNTS OF COLLECTIVE DECISION-MAKING

The minimal liberal account

The minimal liberal account holds that it is (often) sufficient to make collective decisions only on specific actions or policies, while leaving underlying reasons undecided. Inspired by Sunstein (1994), I call such decisions *incompletely theorized* ones.³ The minimal liberal account holds that making collective decisions on more fundamental reasons is (often) not only unnecessary, but also infeasible and undesirable. It is unnecessary because such reasons are practically irrelevant once there is enough agreement on the specific decision made. It is infeasible because people may fail to reach agreement on such reasons. It is undesirable because people may find some reasons offensive and divisive, even when they agree on the resulting decision.

³ Crucially, Sunstein’s own account of ‘incompletely theorized agreements’ refers only to judicial decisions; elsewhere Sunstein is a defender of deliberative democracy with regard to political decisions. Here I use the notion of ‘incomplete theorization’ in a more general sense than intended by Sunstein, so as to apply not only to judicial decisions, but to any kinds of collective decisions.

As noted above, a policy to protect a natural habitat may be widely supported, even when people disagree about *why* the habitat should be protected. If it were decided that it should be protected to respect God's creation, this might offend an atheist environmentalist or someone of another religion, who would otherwise happily endorse protecting the habitat. Similarly, the citizens of a liberal state may agree that human beings have inalienable rights, while disagreeing on where these rights come from. If the state were to take a particular philosophical stance on the source of human rights, this might alienate citizens with a different philosophical view, who endorse human rights for different reasons. Also, if there is broad support for human rights, it is unclear what the state would gain by committing itself to a particular stance on the source of those rights.

In summary, on the minimal liberal account, incompletely theorized agreements are essential for the functioning of a pluralistic society, which is characterized by irresolvable disagreements on 'higher-order' reasons.

The comprehensive deliberative account

In contrast to the minimal liberal account, the comprehensive deliberative account holds that it is (often) insufficient to make collective decisions merely on specific actions or policies, without making the underlying reasons public. The account demands what Pettit (2001b) calls the *collectivization of reason*: It holds that it is (often) necessary for democratic legitimacy to supplement collective decisions on actions or policies with supporting reasons.⁴ These reasons should themselves be collectively decided and publicly defensible. On this account, it is not enough for the legitimacy of an action or policy that the majority endorses this action or policy. Such majority endorsement might stem from a spontaneous majority passion or lack any reasoned justification.

The comprehensive deliberative account proposes two demands to prevent majority tyranny or unreasoned majority decisions and to enhance democratic legitimacy. First, collective decisions should be supported by publicly decided reasons; second, the overall system of collectively decided reasons and outcomes should be coherent. I call the first the *conversability demand*, and the second the *integrity demand*, inspired by Pettit

⁴ Pettit's precise understanding of the collectivization of reason goes further; he thinks of a group as collectivizing reason when it holds itself as a collectivity to the demands of reason.

(2001a) and Dworkin (1986), respectively.⁵ An agent is *conversable* if she can justify her actions by giving appropriate reasons for these actions. An agent displays *integrity* if she is coherent in her judgments or decisions across multiple issues and/or across time.⁶

It is a familiar requirement that individual agents should be conversable and display integrity in this way. To defend their actions when challenged, individuals are usually expected to give reasons for these actions: they are expected to be conversable. Likewise, individuals are usually expected to be rational, which requires coherence of their systems of judgments or decisions: a demand of integrity. The comprehensive deliberative account extends these requirements to groups.

Comparing the two accounts

The minimal liberal and comprehensive deliberative accounts view the relation between reason-giving and the stability and legitimacy of collective decisions very differently. One sees reason-giving as unnecessary for, and potentially threatening to, stability and legitimacy, the other as conducive to, and potentially necessary for, it. To make a collective decision stable and legitimate, on the minimal liberal account, the decision must be *as neutral as possible* between different supporting reasons, as there may be irresolvable disagreements on such reasons. On the comprehensive deliberative account, by contrast, the decision must be *supported* by appropriate reasons, which are also collectively decided and publicly defensible.

The minimal liberal and comprehensive deliberative accounts are simplified ideals, but several accounts in the literature – or variants of such accounts – can be placed at intermediate positions between them. At the risk of oversimplification, here are some examples. With regard to judicial decisions, Sunstein’s account of ‘incompletely theorized agreements’ (1994) is closer to the minimal liberal end of the spectrum, while

⁵ On ‘conversability’, see also Pettit and Smith (1996). Dworkin’s account of ‘integrity’ concerns judicial interpretation, rather than collective decision-making more generally. Here I use the notion of ‘integrity’ in a more general sense than intended by Dworkin.

⁶ Kornhauser and Sager (2004) suggest a taxonomy of integrity constraints. First, an integrity constraint may be applied to the (1a) political or (1b) judicial realms. Second, it may be applied (2a) only to actions or outcomes or (2b) also to supporting reasons and justifications. Third, it may be applied (3a) synchronically or (3b) diachronically. Kornhauser and Sager argue that Pettit’s and Dworkin’s accounts of conversability or integrity occupy opposite positions on this taxonomy, claiming that Pettit is mainly concerned with (1a)-(2b)-(3a) and Dworkin with (1b)-(2a)-(3b). But this is not entirely correct, as Pettit also emphasizes diachronic generalizations of the discursive dilemma and argues that the dilemma occurs in several realms.

Dworkin's account of 'integrity' in law (1996) is closer to the comprehensive deliberative end. With regard to political decisions, Riker's 'liberalism' (as opposed to 'populism') (1982) is closer to the minimal liberal end, while many accounts of deliberative democracy, like Pettit's proposals on the 'collectivization of reason' (2001b), are closer to the comprehensive deliberative end. Rawls's 'political liberalism' (1993) is an important compromise position, on which I make some remarks below. Of course, the ideal-typical accounts discussed here miss out on many aspects and nuances of the actual accounts held by these authors, but a comparison of the ideals still sheds light on those characteristics that actually held accounts share with them.

The difference between the minimal liberal and comprehensive deliberative accounts might seem to lie only in whether or not the reasons behind a collective decision are made public, but not in what the outcome of the decision is. If this were correct, then the choice between the two accounts would perhaps affect people's *perception* of a collective decision – especially the perceived legitimacy, which might be influenced by the public reasons for the decision – but it would *not* affect the decision made. But I argue that the two accounts suggest the use of two different decision procedures that may lead to different outcomes even for the same decision problem.

III. THE DISCURSIVE DILEMMA

Examples

The tenure example. A university committee has to decide whether to give tenure to a junior academic (the *outcome* or *conclusion*). The requirement for tenure is excellence in both teaching and research (the two *reasons* or *premises*). The first among three committee members thinks the candidate is excellent in teaching but not in research; the second thinks she is excellent in research but not in teaching; the third thinks she is excellent in both. So a majority considers the candidate excellent in teaching, a majority considers her excellent in research, but only a minority – the third committee member – thinks the candidate should be given tenure. How should the committee decide?⁷

The global warming example. An expert policy advisory committee has to judge whether, given present greenhouse gas emissions, there will be significant global

⁷ This example is given in Bovens and Rabinowicz (2005).

warming (the *conclusion*). The experts consider two propositions: first, that greenhouse gas emissions are above some critical threshold (the first *premise*); and second, that, if greenhouse gas emissions are above the given threshold, *then* there will be significant global warming (the second *premise*). The two propositions jointly entail that there will be significant global warming. The first among three experts thinks both propositions are true; the second thinks the first but not the second is true; the third thinks the second but not the first is true. So a majority considers each proposition true, yet a majority believes there will *not* be significant global warming. What advice should the committee give?⁸

The habitat protection example. A legislature has to decide whether to protect a natural habitat as a national park. There are many possible reasons for protecting the habitat. Aesthetic reasons: it is a place of great natural beauty. Ecological reasons: its biodiversity is unique. Religious reasons: it is a place of religious significance. Economic reasons: it attracts eco-tourists. Scientific reasons: it provides records of evolutionary processes. Medical reasons: its flora and fauna provide ingredients for new drugs. Political reasons: protecting the habitat boosts the country's reputation as being environmentally friendly. And so on. Each possible reason might serve as a sufficient *premise* to support the *conclusion* to protect the habitat. The protection of the habitat is unanimously supported by the legislators. But the legislature might also seek to agree on a declaration as to *why* the habitat should be protected. Every proposed wording is rejected by a majority. How should the legislators proceed?

Each of these decision problems involves reasons and outcomes, or premises and conclusions. In the first two examples, the relation between premises and conclusion is *conjunctive*: accepting all premises is necessary and sufficient for accepting the conclusion. In the third example, it is *disjunctive*: accepting at least one premise is sufficient (and necessary, if the premises are exhaustive) for accepting the conclusion. Let R_1, R_2, \dots, R_k denote the premises and C the conclusion. In a *conjunctive decision problem*, the relation between premises and conclusion is given by the proposition $C \leftrightarrow (R_1 \ \& \ R_2 \ \& \ \dots \ \& \ R_k)$; in a *disjunctive decision problem*, by $C \leftrightarrow (R_1 \ \text{or} \ R_2 \ \text{or} \ \dots \ \text{or} \ R_k)$.⁹

⁸ This example is given in Dietrich and List (2005); for a related example, see Kornhauser and Sager (2004).

⁹ The first and third examples are direct instances of these conjunctive and disjunctive forms. The second example can be interpreted as an instance of the conjunctive form, but also allows an alternative

Two Decision Procedures

As the minimal liberal and comprehensive deliberative accounts assign different roles to reasons as opposed to outcomes in collective decisions, they also suggest different procedures for solving conjunctive and disjunctive decision problems. On the minimal liberal account, all that is required in each case is a decision on the conclusion. This suggests the use of the *conclusion-based decision procedure*.

The conclusion-based procedure. The group makes a collective judgment on the conclusion (C) by taking a majority vote on that conclusion. The group abstains from making any collective judgments on the premises (R_1, R_2, \dots, R_k).

In this way, judgments on the premises remain private. They influence the collective judgment on the conclusion at most indirectly, through their influence on the individuals' judgments on that conclusion. On the comprehensive deliberative account, by contrast, a decision on the conclusion is not enough. Rather, the decision has to be supported by appropriate reasons, which are themselves publicly decided. This suggests the use of the *premise-based decision procedure*.

The premise-based procedure. The group first makes collective judgments on all premises (R_1, R_2, \dots, R_k) by taking a majority vote on each premise. The group then derives its collective judgment on the conclusion (C) from these collective judgments on the premises, using the appropriate logical relation.¹⁰

In this way, the logical reasoning from the premises to the conclusion is performed at the collective level; the collective judgment on the conclusion is thus consistent with, and supported by, collective judgments on the premises.

The premise- and conclusion-based procedures clearly differ in whether or not any collective judgments are made on the premises; but they may also differ in the collective judgment they generate on the conclusion, as shown next. This problem has become known as the 'doctrinal paradox' (Kornhauser and Sager 1986, 1993) or 'discursive dilemma' (Pettit 2001a) and has sparked a technical literature on judgment aggregation (List and Pettit 2002 and other contributions reviewed above).

representation. The two premises can be written as P and $P \rightarrow Q$, and the conclusion as Q; now the precise logical structure of the problem depends on the interpretation of the conditional \rightarrow . For simplicity, I focus on conjunctive and disjunctive decision problems here. The model of judgment aggregation developed in List and Pettit (2002, 2004) and subsequent papers permits any logical relations expressible in logic.

¹⁰ $C \leftrightarrow (R_1 \& R_2 \& \dots \& R_k)$ in a conjunctive problem; $C \leftrightarrow (R_1 \text{ or } R_2 \text{ or } \dots \text{ or } R_k)$ in a disjunctive one.

Disagreements between the premise- and conclusion-based procedures

Consider the tenure and global warming examples. Both are conjunctive decision problems with premises R_1 and R_2 and conclusion C . The individual judgments are as shown in Table 1. All individuals have consistent judgments relative to the logical relation $C \leftrightarrow (R_1 \& R_2)$.

Table 1: The Discursive Dilemma (Conjunctive Form)

	R_1	R_2	$C \leftrightarrow (R_1 \& R_2)$	C
Individual 1	True	True	True	True
Individual 2	True	False	True	False
Individual 3	False	True	True	False
Majority	True	True	True	False

Under the conclusion-based procedure, the group takes a majority vote on C : here C is rejected. Under the premise-based procedure, the group takes majority votes on R_1 and R_2 , which are both accepted, and deduces its collective judgment on C from its judgments on R_1 and R_2 , using the logical relation $C \leftrightarrow (R_1 \& R_2)$; here C is accepted. So the premise- and conclusion-based procedures produce different outcomes on C . In the tenure example, the candidate would be granted tenure under the premise-based procedure, but not under the conclusion-based one. In the global warming example, the expert committee would advise that there will be significant global warming if it uses the premise-based procedure, but not if it uses the conclusion-based one.

Next consider the habitat example. Let C be the proposition to protect the habitat, and R_1 , R_2 and R_3 the propositions that this should be done, respectively, for scientific, religious and economic reasons. For simplicity, suppose that these are the only relevant reasons, and that there are three legislators (or equal-sized factions) with individual judgments as shown in Table 2. All individuals have consistent judgments relative to the logical relation $C \leftrightarrow (R_1 \text{ or } R_2 \text{ or } R_3)$.

Table 2: The Discursive Dilemma (Disjunctive Form)

	R_1	R_2	R_3	$C \leftrightarrow (R_1 \text{ or } R_2 \text{ or } R_3)$	C
Individual 1	True	False	False	True	True
Individual 2	False	True	False	True	True
Individual 3	False	False	True	True	True
Majority	False	False	False	True	True

Under the conclusion-based procedure, the group takes a majority vote on C : here C is (unanimously) accepted. Under the premise-based procedure, the group takes majority votes on R_1 , R_2 and R_3 , which are each rejected; every proposed reason for protecting the

habitat is rejected by a majority. The collective judgment on C is then deduced from these judgments on R_1 , R_2 and R_3 , using the logical relation $C \leftrightarrow (R_1 \text{ or } R_2 \text{ or } R_3)$; here C is rejected. Again, the premise- and conclusion-based procedures differ in the outcome on C. So the habitat protection policy would be adopted under the conclusion-based procedure, but not under the premise-based one.

The frequency of disagreements between the conclusion- and premise-based procedures

Are the disagreements between the premise- and conclusion-based procedures contrived artifacts, or are they likely to occur in practice? To give a partial answer to this question, let me suggest some considerations on the frequency of such disagreements.

A scenario of likely disagreements. In a disjunctive decision problem, suppose each individual has a small but not very small probability p of accepting each premise.¹¹ Specifically, p is between 0.29 and 0.5 if there are two premises, between 0.21 and 0.5 if there are three and between 0.16 and 0.5 if there are four, generally between $1 - \sqrt[k]{0.5}$ and 0.5. In the habitat example, each proposed reason for protecting the habitat may appeal to a randomly selected individual only with a low, but not very low, probability. Then, with increasing group size, the probability of a disagreement between the premise- and conclusion-based procedures approaches 1. This result holds because, with increasing group size, the probability of a majority on the conclusion approaches 1 (as each individual jointly rejects all premises with a probability less than $(1-p)^k < 0.5$, i.e. accepts at least one premise with a probability greater than 0.5), while the probability of a majority against every premise also approaches 1. A similar result holds in a conjunctive decision problem if each individual has a large but not very large probability p of accepting each premise, specifically between 0.5 and 0.71 if there are two premises, between 0.5 and 0.79 if there are three and between 0.5 and 0.84 if there are four, generally between 0.5 and $\sqrt[k]{0.5}$.

A scenario of unlikely disagreements. In a disjunctive decision problem, suppose at least one premise is more likely than not to be accepted by each individual. For example, one premise might stand out as appealing. Then, with increasing group size,

¹¹ In this and the next scenario, I assume that the judgments of different individuals are mutually independent; also, by saying x is *between* y and z , I mean $y < x < z$.

the probability of a disagreement between the premise- and conclusion-based procedures approaches 0. This result holds because, with increasing group size, the probability that at least one premise is accepted by a majority (and therefore that the conclusion is accepted under both procedures) approaches 1. A similar result holds in a conjunctive decision problem if at least one premise is more likely than not to be rejected by each individual.

The counting and expectation approaches. To quantify the frequency of a given voting paradox, social choice theorists often ask the following two questions. First, given all possible combinations of inputs (here: individual judgments) from a group of individuals, how large is the proportion among those combinations leading to the paradox in question? Call this the *counting approach*.¹² Second, suppose we do not know the probability with which each group member endorses each possible input and we consider every possible probability distribution equally conceivable;¹³ then, what is the expected frequency of the paradox in question? Call this the *expectation approach*.¹⁴ Although answers to these questions are well known for standard voting paradoxes (Gehrlein and Fishburn 1976; Gehrlein 1981, 1983), they have only recently been given for the discursive dilemma, specifically for a disagreement between the premise- and conclusion-based procedures in a two-premise conjunctive decision problem; the answers for a two-premise disjunctive decision problem are the same (List 2005). On the counting approach, the frequency of a disagreement between the two procedures approaches 0.25 with increasing group size; on the expectation approach, it approaches 0.125. Compare this with the frequency of a Condorcet voting cycle over three candidates.¹⁵ On the counting approach, that frequency approaches 0.08774 with increasing group size; on the expectation approach, it approaches 0.0675. Both numbers are smaller than those for disagreements between the premise- and conclusion-based procedures. Although these results do not prove that such disagreements – discursive dilemmas – are empirically likely, they underline that they should be taken as seriously as standard voting paradoxes.

¹² This approach can be shown to be equivalent to the so-called *impartial culture* approach.

¹³ Note that this is a second-order uncertainty. We do not know the judgments that will be made (a first-order uncertainty), *and* we do not know the probability distribution over all possible combinations of judgments (a second-order uncertainty). We apply an epistemological indifference principle at this second-order level and assume each possible probability distribution to be equally likely to be correct.

¹⁴ This approach can be shown to be equivalent to the so-called *impartial anonymous culture* approach.

¹⁵ This comparison is arguably appropriate, as a two-premise conjunctive discursive dilemma is structurally similar to a three-candidate Condorcet paradox (List 2005).

IV. THE PUBLIC AGENDA

The concept of the public agenda

For any given collective decision problem, let the *public agenda* be the set of propositions to be collectively decided. The public agenda might be interpreted as the ‘sphere of public reason’ for that decision problem.

The minimal liberal and comprehensive deliberative accounts can be seen as alternative views on how large or small the public agenda should be. The minimal liberal account recommends a small public agenda: it should include only those propositions on which a decision is needed for practical action. In the habitat protection example, it might include only the proposition to protect the habitat, but no supporting reasons or background considerations. The comprehensive deliberative account recommends a larger public agenda: it should include not only practical propositions, but also relevant supporting reasons and background considerations. In the habitat case, it might include not only the proposition to protect the habitat, but also other considerations – perhaps economic, political, scientific or cultural ones – relevant to that main practical issue.

The minimal liberal and comprehensive deliberative accounts represent extreme positions on how to specify the public agenda. On an intermediate account, some reasons and background considerations might be deemed publicly admissible, others not; accordingly, the public agenda might include the former but not the latter. Rawls’s ‘political liberalism’ (1993), for example, recommends the inclusion of certain reasons and background considerations in the public agenda and the exclusion of others. Those to be included are ‘political’ ones that may be supported by an overlapping consensus in society; those to be excluded are ‘metaphysical’ ones, on which there may be irresolvable disagreement and which do not require public resolution.¹⁶

The public agenda in a conjunctive or disjunctive decision problem

Consider a conjunctive or disjunctive decision problem. Suppose the conclusion *C* is the only practical issue, while the premises R_1, R_2, \dots, R_k are supporting reasons or theoretical considerations. Then the minimal liberal account would include in the public agenda only proposition *C*, while the comprehensive deliberative account would also

¹⁶ An important aspect of Rawls’s ‘political liberalism’ not captured by my analysis is Rawls’s distinction between decisions concerning the basic structure of society and other decisions.

include propositions R_1, R_2, \dots, R_k and a proposition on the logical relation between the premises and the conclusion, such as $C \leftrightarrow (R_1 \& R_2 \& \dots \& R_k)$ or $C \leftrightarrow (R_1 \text{ or } R_2 \text{ or } \dots \text{ or } R_k)$ – or more than one such proposition if the logical relation itself is under debate.

Further, I assume that, when a proposition is included in the public agenda, so is its negation; thus the group can accept either the proposition or its negation or neither.

The complexity of the public agenda

Let me call the public agenda *simple* if it includes only a single ‘atomic’ proposition such as C (or multiple unconnected ones)¹⁷ and *non-simple* if it includes two or more ‘atomic’ propositions such as R_1, R_2 and C and at least one ‘compound’ proposition such as $R_1 \& R_2$ or $C \leftrightarrow (R_1 \& R_2)$ for representing logical relations between other propositions.¹⁸ The minimal liberal account usually recommends a simple public agenda, the comprehensive deliberative account a non-simple one. It is important to emphasize that, generally, the members of a group may disagree not only on premises or conclusions but also on logical relations.

V. JUDGMENT AGGREGATION: A FORMAL MODEL

The discursive dilemma generalized

Suppose the public agenda has been specified, i.e. the propositions to be decided are given. Then the group needs to find a procedure for making those decisions. Let a *decision procedure* be a method of generating collective judgments on the propositions in the public agenda based on the group members’ individual judgments on these propositions.¹⁹ Although I have already introduced two such procedures, i.e. the premise-

¹⁷ And its (their) negation(s).

¹⁸ Crucially, logical interconnections between propositions are *not* taken to be exogenously given. Rather, collective judgments can be made on both ‘atomic’ and ‘compound’ propositions. Technically, an agenda is *non-simple* if it is of one of the following forms (or a superset thereof): (i) it includes $k > 1$ premises R_1, \dots, R_k and either their conjunction $R_1 \& \dots \& R_k$, or their disjunction $R_1 \text{ or } R_2 \text{ or } \dots \text{ or } R_k$, or both (and the negations of all these propositions); (ii) it includes $k > 1$ premises R_1, \dots, R_k , a conclusion C and either the proposition $C \leftrightarrow (R_1 \& \dots \& R_k)$ or the proposition $C \leftrightarrow (R_1 \text{ or } R_2 \text{ or } \dots \text{ or } R_k)$ or both (and negations); (iii) it includes a premise R , a conclusion C and the proposition $R \rightarrow C$ (and negations). The property *non-simple* is slightly more demanding than that of being *not* simple. For generalizations of the property *non-simple* under which proposition 2 below continues to hold, see Dietrich (2005b) and Dietrich and List (2005).

¹⁹ A *decision procedure* is a function whose input is a vector of individual judgment sets and whose output is a collective judgment set (List and Pettit 2002). An individual or collective *judgment set* is a subset of the public agenda interpreted as the set of propositions accepted (as true) by the individual or collective. Individual judgment sets are usually assumed to be *complete* (a member of each proposition-negation pair

and conclusion-based ones, let me take a step back and ask what demands one might impose on a decision procedure. At least three demands seem plausible.

First, the decision procedure should function under conditions of pluralism. Second, the collective judgments generated by the procedure should be responsive to the group members' individual judgments. Third, the collective judgments should display at least some degree of integrity. Unlike the first two demands, the third demand might initially appeal only to proponents of the comprehensive deliberative account. But the integrity demand as stated here is consistent with the minimal liberal account, as it has no force when the public agenda is simple: If the group considers only a single proposition or several unconnected ones, then integrity is vacuously satisfied, no matter how the group decides. Consider the following as a first formalization of the three demands.

Pluralism. The decision procedure accepts as admissible inputs all possible combinations of individual judgments on the propositions in the public agenda.

Majoritarian responsiveness. The collective judgment on each proposition in the public agenda is the judgment of the majority of individuals on that proposition.

Integrity. The collective judgments generated by the decision procedure are consistent, i.e. the propositions collectively accepted *can* be simultaneously true.

While the formalizations of the pluralism and integrity demands are fairly straightforward, majoritarian responsiveness captures only one form of responsiveness and is generalized and relaxed below. My reason for considering it as an initial formalization of responsiveness is that it is the form of responsiveness invoked in discussions of the discursive dilemma. Are there any decision procedures that simultaneously meet the three demands in the present form? If the public agenda is non-simple, as in standard conjunctive or disjunctive decision problems, the answer is negative.

Proposition 1. For a non-simple public agenda, there exists no decision procedure that meets the demands of pluralism, majoritarian responsiveness and integrity.²⁰

is accepted) and *consistent* (in the standard sense of logic), but these assumptions can be relaxed. In the collective case, completeness (decisiveness) and consistency (integrity) are introduced as demands on decision procedures.

²⁰ Assuming three or more individuals. For exactly four individuals, note the remark after Table A3. A variant of the proposition holds whenever the largest minimal inconsistent subset of the public agenda has three or more elements.

To sketch a proof, assume, for a contradiction, that there exists a decision procedure satisfying pluralism, majoritarian responsiveness and integrity. For any group of three or more individuals, combinations of individual judgments similar to those in Tables 1 and 2 can be constructed; an explicit construction is given in the Appendix. By the demand of pluralism, those combinations are admissible inputs to the decision procedure. By the demand of majoritarian responsiveness, the collective judgment on each proposition is the majority judgment on that proposition. But, for the given combinations of individual judgments, those majority judgments are not consistent: the accepted propositions – for example, R_1 , R_2 , $C \leftrightarrow (R_1 \ \& \ R_2)$ and not- C in Table 1 – cannot be simultaneously true. This contradicts the integrity demand.

A more general impossibility result

I now keep the formalizations of the pluralism and integrity demands, but generalize that of the responsiveness demand. To give an alternative formalization of responsiveness without the majoritarian requirement, define a *single-proposition decision rule* to be a method of summarizing the individual judgments on a single proposition into a collective judgment on that proposition. Examples are not only simple majority voting, but also qualified majority voting, sub-majority voting, the unanimity rule, a dictatorship of one individual or even quirkier rules that decide a proposition based on whether it is supported by an even or odd number of individuals.

Systematic responsiveness. There exists a single-proposition decision rule such that the collective judgment on each proposition in the public agenda is the result of applying that decision rule to the individual judgments on that proposition.

Although weaker than majoritarian responsiveness, systematic responsiveness is still a strong requirement, as it demands independent and equal treatment of all propositions. Below I consider cases where we may have good reasons to abandon the requirement of independent and equal treatment across propositions and thus to relax systematic responsiveness. In addition to replacing majoritarian responsiveness with systematic responsiveness here, let me also make explicit two demands implicit in majoritarian responsiveness: first, the decision procedure should give all individuals

equal weight; and second, the procedure should not leave any propositions in the public agenda undecided (after all, those propositions are precisely the ones to be decided):²¹

Anonymity. The decision procedure is invariant under permutations of the individuals.

Decisiveness. The decision procedure generates a determinate collective judgment on every proposition in the public agenda (acceptance of either the proposition or its negation).

Again, for a non-simple public agenda, there exists no decision procedure that meets all demands simultaneously, as proved by List and Pettit (2002).

Proposition 2. For a non-simple public agenda, there exists no decision procedure that meets the demands of pluralism, systematic responsiveness, anonymity, decisiveness and integrity.²²

This impossibility result is analogous, but not equivalent, to Arrow's impossibility theorem (1951) (for a comparison, see List and Pettit 2004; Dietrich and List 2005). If the public agenda is not only non-simple, but includes more complex logical interconnections (which can be technically described), then an impossibility result still occurs if systematic responsiveness is weakened to *independence*,²³ but *respect for unanimity*²⁴ is demanded (Pauly and van Hees 2005; Dietrich 2005; Dietrich and List 2005; for further results, see the references in an earlier note).

Escape-routes

Propositions 1 and 2 show that some plausible demands on a decision procedure are mutually incompatible when the public agenda is non-simple. As the comprehensive deliberative account typically recommends a non-simple public agenda, the results pose a challenge for that account. *If* the account requires all the identified demands, *then* it faces an impossibility problem, as the demands cannot be simultaneously met. For the minimal

²¹ These demands are implicit in majoritarian responsiveness. All individuals have equal weight in majority decisions, and majority decisions (almost) always lead to a determinate judgment (except for majority ties).

²² Again, assuming three or more individuals.

²³ Independence is the demand that, for each proposition in the public agenda, there exists a single-proposition decision rule such that the collective judgment on that proposition is the result of applying that decision rule to the individual judgments on that proposition. (Unlike systematic responsiveness, independence allows the single-proposition decision rule to vary from proposition to proposition.)

²⁴ Defined below.

liberal account, the problem is less serious, as that account requires only a simple public agenda. How can the impossibility problem be avoided? Logically, there are several escape-routes (List and Pettit 2002), but only some are appealing from minimal liberal or comprehensive deliberative perspectives.

Shrinking the public agenda. If the public agenda is simple – containing only a single proposition or multiple *unconnected* ones – then the group can take majority votes on those (few) propositions in the public agenda and will meet all the demands introduced above. This decision procedure is precisely the conclusion-based procedure in case the premises in a conjunctive or disjunctive decision problem are *excluded* from the public agenda, as recommended by the minimal liberal account. So the present way to avoid the impossibility problem is the preferred one from a minimal liberal perspective.

Reducing pluralism, especially through deliberation. If not all possible combinations of individual judgments are admissible as inputs to the decision procedure, but only ones with sufficient agreement among individuals, then the group can take majority votes on all propositions in the public agenda – even a non-simple one – and will meet all the demands imposed above, except that of pluralism. For a certain type of structural agreement among individuals, called *unidimensional alignment*, majoritarian responsiveness and integrity are compatible (List 2003). This escape-route is of interest from a comprehensive deliberative perspective. Proponents of that account will typically advocate not only public decision-making on reasons, but also public deliberation in a decision process. Such deliberation might transform individual judgments and generate the required structural agreement among group members (Miller 1992). There is some evidence for a mechanism along these lines in a preference (as opposed to judgment) aggregation context, but, depending on the original level of disagreement in a group, such a mechanism will not always be operational (List, Fishkin, Luskin and McLean 2000).

Relaxing the responsiveness demand. Instead of requiring that the collective judgment on *every* proposition in the public agenda should be the majority judgment on that proposition (or the result of applying the same single-proposition decision rule), one might require that collective judgments should be responsive to individual ones in this way only *for some but not all* propositions, and that collective judgments on others should be derived indirectly. The premise-based procedure is an instance of this. Here the

collective judgments on the premises are the majority judgments on these premises, but the collective judgment on the conclusion is derived from these judgments using the appropriate logical relation. This approach makes explicit an *order of priority* among the propositions: some serve as reasons or premises for others and are given priority in a (collective) reasoning process. This idea is generalized in Section VIII below. Proponents of the comprehensive deliberative account might argue that the ‘collectivization of reason’ is captured more adequately by this approach than by pure proposition-by-proposition aggregation (Pettit 2001a; Chapman 2002). On this view, relaxing majoritarian or systematic responsiveness is seen not as a sacrifice but as a gain.

Appointing a dictator. The demands of pluralism, systematic (though not majoritarian) responsiveness, decisiveness and integrity are met for a non-simple public agenda if and only if one individual is appointed as a dictator who determines the group’s judgments (Pauly and van Hees 2005; Dietrich and List 2005). This decision procedure is unappealing from both minimal liberal and comprehensive deliberative perspectives. It not only violates anonymity, a fundamental liberal principle, but also abandons the idea that collective judgments should be the result of some non-dictatorial collective process.

Relaxing the decisiveness demand. Instead of making a determinate collective judgment on *every* proposition in the public agenda, the group might abstain from making judgments on certain (controversial) propositions, such as when there is not enough agreement for or against a proposition. Specifically, the group might apply a supermajority or unanimity rule whereby (i) a proposition is collectively accepted if it is accepted by a large supermajority (or all) of the individuals, (ii) its negation is collectively accepted if that negation is accepted by a similarly large supermajority, (iii) otherwise the proposition is left undecided. If the required supermajority is sufficiently large, this decision procedure will meet all the demands introduced above, except that of decisiveness, as many propositions will be left undecided (List 2004; Goodin and List 2006). Compare this approach with that of shrinking the public agenda. On the present approach, the public agenda may remain large (particularly: non-simple), but the group is not required to settle all issues in that agenda. On the approach of shrinking the public agenda, the group *is* required to settle all issues in the public agenda, but, to make this feasible, the agenda is made small. The present approach is of interest from both minimal

liberal and comprehensive deliberative perspectives. Proponents of the minimal liberal account may be attracted to the idea that collective judgments are made only if they are *either* practically necessary *or* reducible to the unanimity (or near-unanimity) judgments of individuals. Compare Buchanan and Tullock's unanimitarian approach (1962) to decision-making on constitutional matters. Proponents of the comprehensive deliberative account may be attracted to the idea that the public agenda can remain large, and perhaps include reasons, even if not all propositions in that agenda are collectively decided.

Relaxing the integrity demand. Finally, a way to avoid the identified impossibility problem is to relax the integrity demand and to allow inconsistent collective judgments. But this route seems unattractive. Proponents of the comprehensive deliberative account will be unwilling to relax the integrity demand, as it lies at the heart of their account. Proponents of the minimal liberal account may be prepared to relax it, but do not need to, as their account largely excludes higher-order reasons and logical relations from the public agenda, so that integrity is an (almost) vacuous requirement.

It is now clear how the premise- and conclusion-based procedures fit into the logical space mapped out by the demands on decision procedures introduced above. Each procedure represents one escape-route from the identified impossibility problem. The conclusion-based procedure represents the route of shrinking the public agenda and permitting an 'incompletely theorized' decision on the conclusion. The premise-based procedure represents that of relaxing responsiveness so as to allow the prioritization of some propositions over others and permit the 'collectivization of reason'. As these two routes are the preferred ones from minimal liberal and comprehensive deliberative perspectives, respectively, the present observations reinforce my earlier claim that the minimal liberal account suggests a conclusion-based approach to collective decision-making, while the comprehensive deliberative account suggests a premise-based one. But, as I argue below, even on the comprehensive deliberative account, a conclusion-based approach is sometimes preferable on epistemological grounds.

VI. THE TRUTH-TRACKING PROPERTIES OF THE PREMISE- AND CONCLUSION-BASED PROCEDURES

How good are the premise- and conclusion-based procedures at tracking the truth in decisions involving matters of fact?

Truth-tracking and reliability

Consider a proposition, P, which is factually true or false, such as the proposition that a patient has some disease. Inspired by Nozick (1981), let me say that an agent (or decision procedure) *tracks the truth* on P if two conditions hold:²⁵ First, whenever P is true, the agent (procedure) judges P to be true. Second, whenever P is false, it judges P to be false.

Most agents (and decision procedures) are fallible and meet these two conditions at most approximately. To quantify an agent's (or procedure's) reliability at truth-tracking, consider two conditional probabilities, corresponding to the two conditions above: first, the conditional probability that the agent (procedure) judges P to be true, given that P is true; second, the conditional probability that it judges P to be false, given that P is false. I call these two conditional probabilities the agent's (procedure's) *positive* and *negative reliability* on P. Sometimes the two probabilities are the same (*symmetrical* reliability), sometimes they come apart: an agent (procedure) may be better at detecting the truth of P than its falsehood, or vice-versa (*asymmetrical* reliability). Many medical diagnostic tests are asymmetrical in this sense: they are better at detecting the presence of some disease if the patient has the disease than its absence if the patient does not.

An agent (or decision procedure) with a high positive *and* negative reliability on P is obviously a good truth-tracker on P, and one with a low positive *and* negative reliability a bad one. But even an agent (procedure) with a high positive and low negative, or low positive and high negative, reliability can be a bad truth-tracker. Imagine a diagnostic test that *always* produces a negative verdict, regardless of whether or not the patient has the disease. This test has a negative reliability of 1, but a positive reliability of 0. As a truth-tracker, it is useless.²⁶

²⁵ To obtain Nozick's analysis, the present two conditionals must be read as counterfactual conditionals and not just as material conditionals. My simple operationalization of truth-tracking in terms of conditional probabilities is arguably weaker than Nozick's counterfactual analysis.

²⁶ This example also shows that it is not always helpful to summarize the positive and negative reliability levels into one figure, by considering the unconditional probability that the agent's judgment matches the

Majority voting and truth-tracking on a single proposition

Suppose a group of individuals – a jury, for example – has to make a collective judgment on some factual proposition, based on the group members’ individual judgments on that proposition. What decision procedure should the group use? The Condorcet jury theorem shows that, under mild assumptions about individual reliability, majority voting is a remarkably good truth-tracker (e.g. Grofman, Owen and Feld 1983).

Condorcet jury theorem. Suppose each group member has a (positive and negative) reliability p greater than 0.5 on some proposition and different individuals’ judgments are mutually independent. Then the (positive and negative) reliability of the group’s majority judgment on that proposition approaches 1 as the group size increases.

So if each group member has a chance greater than 0.5 of making a correct judgment on some proposition (both when it is true and when it is false) – call this *Condorcet’s competence assumption* – then the group can make a fairly reliable judgment on that proposition just by taking a simple majority vote, provided the number of individuals is sufficiently large and their judgments are mutually independent.

Sometimes a group is particularly concerned with avoiding false positive decisions (the acceptance of a false proposition) or avoiding false negative ones (the rejection of a true proposition). In such cases, super- or sub-majority voting instead of simple majority voting is optimal. For example, supermajority voting is often used in criminal trials, to avoid convicting the innocent. Sub-majority voting is sometimes used in legislative or judicial decisions on whether to consider an initiative or complaint, to avoid overlooking valid initiatives or complaints. But without an asymmetry between false positives and false negatives, the Condorcet jury theorem supports the use of simple majority voting in decisions where Condorcet’s competence assumption is met.

truth-value of some proposition P . This unconditional probability is the sum of (i) the unconditional probability that P is true multiplied by the conditional probability that the agent judges P to be true, given that P is true and (ii) the unconditional probability that P is false multiplied by the conditional probability that the agent judges P to be false, given that P is false. Bovens and Rabinowicz (2005) use such a summary figure to measure reliability. To see that this does not always provide a good measure of reliability, note that, even for the indiscriminate diagnostic test of the example, the summary figure can be arbitrarily close to 1, namely when the disease in question is very unlikely. But the fact that (say) CJD is an extremely unlikely disease does not make a test that always says ‘negative’ a good truth-tracker.

The case of multiple propositions

Suppose now a group has to decide multiple propositions, such as premises and conclusions in a conjunctive or disjunctive decision problem, and suppose further the propositions are factual ones. Doesn't the Condorcet jury theorem simply suggest that the group should take majority votes on all those propositions?

If all group members meet Condorcet's competence assumption on *every* proposition, this is true. But, even on an optimistic view, individuals will not generally meet this assumption on multiple interconnected propositions. To see why not, suppose an individual has a (positive and negative) reliability of 0.65 on each of propositions R_1 and R_2 (two factual premises, say) and her judgments on R_1 and R_2 are mutually independent. The probability of 0.65 is well above the threshold of 0.5 in Condorcet's competence assumption. But now consider the same individual's reliability on a conclusion C that is true just in case the conjunction $R_1 \& R_2$ is true. If R_1 and R_2 are both true, the individual will have a probability of $0.65^2 = 0.4225$ of judging $R_1 \& R_2$ (and hence C) to be true. So the individual's *positive* reliability on C is below the threshold of 0.5. Although the individual's *negative* reliability on C is still above that threshold, the individual does not meet Condorcet's competence assumption for the conclusion C , even though she meets that assumption for each of the two premises R_1 and R_2 . If all group members are like this, then majority voting on the conclusion C is a bad truth-tracker – specifically, its positive reliability on C will approach 0 as the group size increases.

In general, individuals have different levels of reliability on different propositions. Some propositions are harder to judge than others. For example, it is harder to detect the truth of a conjunction than that of each conjunct, but easier to detect its falsehood; it is easier to detect the truth of a disjunction than that of each disjunct, but harder to detect its falsehood. Depending on the logical and epistemological nature of a decision problem involving multiple propositions, individuals may have different levels of reliability across these propositions. This needs to be taken into account when a procedure is chosen for making collective judgments on those propositions.

Extending recent work by Bovens and Rabinowicz (2005), I suggest that the premise-based procedure is a superior truth-tracker for some decision problems and the conclusion-based procedure for others (see also List 2005). The conditions under which

the premise-based procedure outperforms the conclusion-based one are similar to those assumed by Bovens and Rabinowicz (although they suggest, debatably,²⁷ that even under these conditions the conclusion-based procedure is sometimes better). The conditions under which the conclusion-based procedure outperforms the premise-based one are explored here for the first time. In my discussion, I keep the standard assumption of the Condorcet jury theorem that the judgments of different individuals are mutually independent (for discussions of dependencies between different individuals' judgments, see Ladha 1992, Estlund 1994, Dietrich and List 2004b).

Conditions under which the premise-based procedure outperforms the conclusion-based procedure

Consider a conjunctive or disjunctive decision problem with k premises R_1, \dots, R_k and a conclusion C .

The structure of individual reliability. Each individual has a (positive and negative) reliability p above 0.5 and below $\sqrt[k]{0.5}$ on each premise, and her judgments on different premises are mutually independent. So each individual meets Condorcet's competence assumption on each premise, but is still fallible.

An example. In the tenure example, the committee members' judgments on the two tenure requirements might satisfy these assumptions if we assume, first, that each committee member is reasonably, but not very, good at judging the candidate's teaching performance and at judging her research performance, and, second (more idealistically), that each member judges the first tenure requirement independently from the second.

The implication for a conjunctive decision problem. Under the premise-based procedure, the group's (positive and negative) reliability on the conclusion approaches 1 with increasing group size. In the tenure example, a large committee is likely to grant the candidate tenure if and only if she meets the requirements. By contrast, under the conclusion-based procedure, the group's negative reliability approaches 1, but its positive reliability approaches 0. Here the conclusion is likely to be rejected in a large group, even when it is true. In the example, the candidate is likely to be denied tenure, no matter how

²⁷ See my earlier note. But I agree that the conclusion-based procedure is sometimes preferable when there are certain asymmetries between false positives and false negatives, as discussed below.

excellent she is. The conclusion-based procedure is similar to a diagnostic test that (almost) always gives the same verdict (here the rejection of the conclusion).

The implication for a disjunctive decision problem. Under the premise-based procedure, the group's (positive and negative) reliability on the conclusion approaches 1 with increasing group size. Under the conclusion-based procedure, its positive reliability approaches 1, but its negative reliability approaches 0. Here the conclusion is likely to be accepted in a large group, even when it is false. For a sketch proof, see the Appendix.

Conditions under which the conclusion-based procedure outperforms the premise-based procedure

First consider a disjunctive decision problem with k premises R_1, \dots, R_k and a conclusion C .

The structure of individual reliability. Suppose each premise is thought by the individuals to have a low prior probability of being true. If *all* premises are false, then each individual has a probability p greater than 0.5 of judging that they are all false. If (at least) one premise is true, then each individual has a probability p greater than 0.5 of judging one to be true, but it is still difficult for the individual to determine which one it is. The individual's probability p of judging one premise to be true is distributed across the k premises: she has a probability p_1 of judging R_1 true, a probability p_2 of judging R_2 true, and so on, where $p = p_1 + \dots + p_k$ exceeds 0.5, but each of p_1, \dots, p_k is less than 0.5. (One might also assume that, if R_i is the premise that is actually true, then p_i exceeds p_j for any j other than i , but this assumption is not essential.)

An example. Again, an expert committee has to give advice on whether there will be significant global warming. Call this proposition C . Here C is true if and only if *at least one* of several possible causes for global warming, R_1, \dots, R_k , is in operation. The following scenario is conceivable. If there is no global warming – none of the causes is in operation – then each scientist has a chance above 0.5 of ruling out global warming and *all* of its causes. But if there is global warming – at least one cause is in operation – then each scientist has a chance above 0.5 of judging this to be the case, but only a chance below 0.5 of identifying the correct cause. (Perhaps each scientist is still more likely to pick the correct cause than any incorrect one, but none of these probabilities exceeds 0.5.)

The implication. Under the conclusion-based procedure, the group's (positive and negative) reliability on the conclusion approaches 1 with increasing group size. In the new global warming example, a large committee is likely to track the truth on the presence or absence of global warming. Under the premise-based procedure, the group's negative reliability approaches 1, but its positive reliability approaches 0. Here the conclusion is likely to be rejected in a large group, even when it is true. In the example, even when there is global warming, none of the causes is likely to be accepted by a majority of scientists, and hence the committee is likely to advise that there will not be significant global warming. For a sketch proof, see the Appendix.

Next consider a conjunctive decision problem with k premises R_1, \dots, R_k and a conclusion C .

The structure of individual reliability. Suppose each premise is thought by the individuals to have a high prior probability of being true. If *all* premises are true, then each individual has a probability p greater than 0.5 of judging that they are all true. If (at least) one premise is false, then each individual has a probability p greater than 0.5 of judging one to be false, but again it is difficult for the individual to determine which one it is. The individual has a probability p_1 of judging R_1 false, a probability p_2 of judging R_2 false, and so on, where $p = p_1 + \dots + p_k$ exceeds 0.5, but each of p_1, \dots, p_k is less than 0.5. (One might assume that, if R_i is the premise that is actually false, then p_i exceeds p_j for any j other than i , but again this assumption is not essential.)

An example. An expert committee (e.g. the board of a space agency) has to give advice on whether a complex technological system (e.g. a space shuttle) is free from a significant safety risk. This proposition, C , is true if and only if the system's multiple components *all* function correctly. Here the following scenario is conceivable. If the system is free from any significant risk – all components function correctly – then each expert may have a chance above 0.5 of judging this to be the case. But if there is a significant risk – at least one component is faulty – then each expert may have a chance above 0.5 of judging that there is a risk, but only a chance below 0.5 of correctly identifying the faulty component. (Perhaps each expert is more likely to identify the faulty component than any other component, but none of these probabilities exceeds 0.5.)

The implication. Under the conclusion-based procedure, the group's (positive and negative) reliability on the conclusion approaches 1 with increasing group size. In the risk assessment example, a large committee is likely to track the truth on whether or not the system is free from a significant risk. Under the premise-based procedure, the group's negative reliability approaches 0, but its positive reliability approaches 1. Here the conclusion is likely to be accepted in a large group, even when it is false. In the example, even when there is a significant risk, none of the system's components is likely to be identified as faulty by a majority of committee members, and hence the committee is likely to advise that there is no significant risk. For a sketch proof, see the Appendix.

Comparing the conditions

I have argued that, when individuals make judgments on multiple propositions, such as premises and conclusions, it is plausible that they satisfy Condorcet's competence assumption – a greater than 0.5 chance of making a correct judgment – on some propositions but not on others.

Sometimes the truth of the premises is directly discernible, while discerning the truth of the conclusion – e.g. a complex conjunction – requires complicated inferences. In these cases, individuals are more likely to satisfy Condorcet's competence assumption on the premises than on the conclusion. Other times the truth of the conclusion is easier to discern than that of the premises; the conclusion might be a broad, 'coarse-grained' proposition, while the premises might be 'fine-grained' propositions about complex causal mechanisms. In those cases, individuals are more likely to satisfy Condorcet's competence assumption on the conclusion than on the premises.

The key to choosing a decision procedure that tracks the truth on multiple propositions lies in identifying those propositions on which Condorcet's competence assumption is met by the relevant individuals. If the assumption is met on the premises but not on the conclusion, then the premise-based procedure outperforms the conclusion-based one at truth-tracking. If it is met on the conclusion but not on the premises, then the conclusion-based procedure outperforms the premise-based one. Even when the public agenda is defined broadly, as recommended by the comprehensive deliberative account, the premise-based procedure is sometimes a worse truth-tracker than the conclusion-

based one. Therefore even the proponents of the comprehensive deliberative account may sometimes opt for a conclusion-based approach if they are concerned with truth-tracking.

False positives versus false negatives

As noted, in some decision problems there is an asymmetry between false positives and false negatives. In the tenure example, on an employee-friendly policy, it is worse not to grant tenure to a deserving candidate than to grant tenure to an undeserving one; on an employer-friendly policy, the asymmetry is reversed. In the global warming example, it is worse from a precautionary perspective not to take action (e.g. reduce greenhouse gas emissions) when global warming is a serious threat than to take action when the threat is less serious than feared. In the habitat example, it is worse from a conservationist perspective not to protect a significant habitat than to protect a less significant one.

In a conjunctive decision problem, the conditions under which the conclusion-based procedure accepts the conclusion are more demanding than those under which the premise-based procedure does. The conclusion-based procedure accepts the conclusion only when a majority of individuals jointly accepts all premises, while the premise-based procedure accepts it already when separate (and possibly non-overlapping) majorities accept each premise without necessarily accepting the conclusion. So the conclusion-based procedure produces fewer false positives than the premise-based one, but at the expense of producing more false negatives.

In a disjunctive decision problem, the reverse is true. The conditions under which the premise-based procedure accepts the conclusion are more demanding than those under which the conclusion-based procedure does. The premise-based procedure accepts the conclusion only when at least one premise is accepted by a majority of individuals, while the conclusion-based procedure accepts it already when a majority of individuals accepts the conclusion based on different premises (which may each be accepted only by a minority). Here the premise-based procedure produces fewer false positives than the conclusion-based procedure, but at the expense of producing more false negatives.

This suggests that, when the underlying structure of a decision problem is conjunctive, a minimal liberal approach is less likely to produce false positive decisions,

but more likely to produce false negative ones, than a comprehensive deliberative approach; the reverse is true when the underlying structure of a problem is disjunctive.

Based on the asymmetries in the examples above, one might therefore recommend a minimal liberal approach in the habitat case and a comprehensive deliberative one in the global warming case. In the tenure case, a minimal liberal approach is preferable from the perspective of a university concerned with protecting its reputation, whereas a comprehensive deliberative one is preferable from the perspective of a junior faculty member applying for tenure.

VII. THE STRATEGIC INCENTIVES GENERATED BY THE PREMISE- AND CONCLUSION-BASED PROCEDURES

Mechanical versus behavioural differences of decision procedures

Different decision procedures are often not only *mechanically* different, in that they generate different collective outputs (decisions) for the same individual inputs (votes or judgments); they are often also *behaviourally* different, in that they lead decision-makers to adjust their behaviour differently and thereby affect the inputs decision-makers feed into these procedures in the first place. This point has been documented extensively in comparisons of different electoral systems (e.g. Norris 2004), and there is every reason to think that it may also be relevant to the issues discussed here.

In my comparison of the decision procedures suggested by the minimal liberal and comprehensive deliberative accounts, I have so far focused on their mechanical differences, holding all behavioural assumptions fixed. Specifically, I have assumed that individuals reveal the same individual judgments independently of whether the setting is a minimal liberal or comprehensive deliberative one – technically, independently of whether the decision procedure is a premise-based or conclusion-based one. On the basis of technical work by Dietrich and List (2004a), I now question these behavioural assumptions and discuss behavioural differences between the two approaches.

An example

Let me revisit the tenure-committee example above, with individual judgments as shown in Table 1. When I argued that the premise-based procedure leads to a positive tenure decision, while the conclusion-based procedure leads to a negative one, I assumed that,

under both procedures, the individuals reveal exactly the judgments shown in Table 1; call these the individuals' *truthful judgments*. Is this assumption plausible?

Table 1 (restated)

	R_1 : The candidate is excellent in teaching.	R_2 : The candidate is excellent in research.	$C \leftrightarrow (R_1 \& R_2)$	C : The candidate deserves tenure.
Individual 1	True	True	True	True
Individual 2	True	False	True	False
Individual 3	False	True	True	False
Majority	True	True	True	False

Suppose that each committee member is convinced of his or her own judgments and has a desire to bring about collective judgments that coincide with those judgments. In particular, suppose, for the moment, that each committee member cares first and foremost about bringing about a tenure decision that coincides with his or her own judgment on what the tenure decision should be: so individuals 2 and 3 would like to bring about a negative tenure decision (of which they are each convinced), while only individual 1 would like to bring about a positive one (of which she is equally convinced). (Below I describe this particular motivation of the committee members as *outcome-oriented* and contrast it with an alternative, *reason-oriented* motivation. Although the present example assumes an outcome-oriented motivation, the more general results on strategic manipulability and strategy-proofness below require no such assumption.)

Now suppose that the committee uses the premise-based procedure. If all individuals reveal their judgments truthfully, then, as shown above, the conclusion will be accepted and the candidate will be granted tenure. Individual 1 wishes to achieve this outcome, but individuals 2 and 3 would be unhappy with it. On reflection, individual 2 realizes that, if she falsely expresses the judgment that the candidate is *not* excellent in teaching (i.e. that R_1 does not hold), contrary to her truthful view, then the majority judgment on the candidate's teaching performance will change from a positive to a negative one (i.e. from the acceptance of R_1 to its rejection); by implication the candidate will be denied tenure, which is individual 2's preferred outcome. Likewise, individual 3 realizes that, if she falsely expresses the judgment that the candidate is *not* excellent in research (i.e. that R_2 does not hold), contrary to her truthful view, she can change the majority judgment on the candidate's research performance and bring about a negative tenure decision, which is individual 3's preferred outcome. So individuals 2 and 3 each

have an incentive to strategically misrepresent their judgments, so as to bring about an outcome they each prefer: a negative tenure decision.

By contrast, suppose the committee uses the conclusion-based procedure. If all individuals reveal their judgments truthfully – here only judgments on the conclusion are revealed – then, as shown above, the conclusion will be rejected and the candidate will be denied tenure. Individuals 2 and 3 wish to achieve this outcome and have every incentive to reveal their negative judgment on the conclusion truthfully. Individual 1 wishes to grant the candidate tenure, but can do no better than to express this judgment truthfully.

In this example, the individuals all have an incentive to reveal their judgments truthfully under the conclusion-based procedure but not under the premise-based one. Curiously, if the individuals all act on their incentives here – which is what rational choice theory predicts – then the outcome under the two procedures is exactly the same. Under both procedures, the candidate is denied tenure, albeit via a different route. Under the conclusion-based procedure, the outcome is brought about through a truthful majority judgment on the conclusion itself. Under the premise-based procedure, it is brought about through strategically manipulated majority judgments on the premises.

Strategic manipulability and strategy-proofness

I have illustrated that the premise-based procedure may give individuals an incentive to strategically misrepresent their judgments, while the conclusion-based procedure, at least in the example, gives individuals an incentive to reveal their judgments truthfully. This suggests that the premise- and conclusion-based procedures may affect the behaviour of individuals differently.

I now address the incentives generated by a decision procedure more generally. Consider any decision procedure, i.e. any method of generating collective judgments on the propositions in the public agenda on the basis of the group members' individual judgments. When does an individual have the *opportunity* to 'manipulate' the collective decision by strategically misrepresenting her judgments?

Manipulability. Given a combination of individual judgments, a decision procedure is *manipulable* by some individual on some proposition if the following two conditions hold (Dietrich and List 2004a): (i) if the individual reveals her judgments on

all propositions truthfully, then the collective judgment on the proposition in question differs from the individual's own judgment on that proposition; (ii) if the individual strategically misrepresents her judgments on some propositions, then the collective judgment on the proposition in question coincides with the individual's truthful judgment on that proposition.

For example, given the combination of individual judgments in Table 1, the premise-based procedure is manipulable by individual 2 on proposition R_1 and manipulable by individual 3 on proposition R_2 . In each case, conditions (i) and (ii) in the definition of manipulability are met.

Now, if an individual has the *opportunity* to manipulate the collective judgment on some proposition (as captured by the definition of manipulability just given) *and* she cares sufficiently strongly about her own judgment on that proposition, then her *opportunity* to manipulate translates into an *incentive* to manipulate, and she may well act on that incentive. A decision procedure that never creates any opportunities for strategic manipulation meets the demand of *strategy-proofness* (Dietrich and List 2004a):

Strategy-proofness. There exists no admissible combination of individual judgments for which the decision procedure is manipulable by any individual on any proposition.

An impossibility result

Are there any strategy-proof decision procedures? More precisely, are there any decision procedures that meet the demand of strategy-proofness together with some of the demands introduced earlier? Clearly, not all of the earlier demands can be imposed, as I have already shown them to be mutually incompatible. In particular, instead of majoritarian or systematic responsiveness, let me require only the following demand, together with the demands of anonymity and decisiveness.

Respect for unanimity. If all individuals unanimously judge some proposition to be true, then the proposition is also collectively judged to be true.

Note that this demand has no force when there is no unanimous agreement between individuals. The following result has been proved by Dietrich and List (2004a).

Proposition 3. For the public agenda of a conjunctive or disjunctive decision problem, there exists no decision procedure that meets the demands of pluralism, strategy-proofness, respect for unanimity, anonymity, decisiveness and integrity.²⁸

Again, the result describes a trade-off between some arguably plausible demands. How can this trade-off be solved? As before, each demand can in principle be relaxed. Let me revisit the preferred escape-routes from the perspectives of the minimal liberal and comprehensive deliberative accounts.

Shrinking the public agenda. If the public agenda is shrunk so as to contain only the conclusion (or multiple *unconnected* propositions), then the impossibility problem disappears. In particular, the conclusion-based procedure is strategy-proof. As a vote is taken only on a single proposition, namely the conclusion, each individual's best strategy – the one most likely to lead to the individual's preferred outcome – is always to reveal her judgment on that conclusion truthfully. Apart from shrinking the public agenda, the conclusion-based procedure meets all other demands used in proposition 3. So the minimal liberal approach provides a way to avoid the present impossibility problem too.

Relaxing responsiveness. Since the responsiveness condition has already been considerably relaxed – neither majoritarian nor systematic responsiveness is required – there is less room for any relaxations here. One might consider relaxing respect for unanimity, although technical work suggests that relaxing this demand alone does not open up promising possibilities.²⁹ Again, relaxing anonymity is unattractive. Finally, relaxing decisiveness is an option just as before, but unanimitarian or supermajoritarian decision procedures – which are strategy-proof – have the potential disadvantage of leaving many propositions undecided.

In summary, while the minimal liberal account has a preferred escape-route from the present impossibility problem at its disposal, the comprehensive deliberative account does not and is therefore faced with a serious challenge here. Ultimately, proponents of the comprehensive deliberative account may have no choice but to relax the formal demand of strategy-proofness itself and *either* to accept the occurrence of strategic

²⁸ Again, assuming three or more individuals. In fact, the proposition holds for a larger class of public agendas called *path-connected* (Dietrich and List 2005a).

²⁹ When the public agenda contains 'rich' logical interconnections, even a weaker demand than respect for unanimity – *weak responsiveness* – is inconsistent with the other demands (Dietrich and List 2004a).

manipulation – which seems undesirable from a comprehensive deliberative perspective – *or* to find other ways of inducing individuals to reveal their judgments truthfully.

As noted above, proponents of the comprehensive deliberative account will advocate not only the publicity of reasons, but also the use of public deliberation in collective decisions. It has been suggested that, when people reveal their views in public deliberation rather than anonymous aggregation, they are less likely to act strategically. There are at least two possible mechanisms for this, an incentive-based one and a psychological one (Mackie 1998; Dryzek and List 2003). According to the incentive-based mechanism, a potential manipulator is at greater risk of being exposed in a setting of social interaction than in an anonymous ballot, and being exposed as a manipulator is costly, both in terms of achieving the preferred decision outcome and in terms of one's reputation (Brennan and Pettit 2004). If this mechanism is in operation in a public deliberative setting, then individuals have less of an incentive to act strategically there. According to the psychological mechanism, people adopt a more cooperative disposition when they socially interact with each other and they focus more on the group's perspective than on their own individual perspective. If this mechanism is in operation in a public deliberative setting, then, again, individuals are less likely to act strategically there, even if strategic incentives are present.

Of course, it is still an open question to what extent any such mechanism can be relied upon to induce individuals to reveal their judgments truthfully. At this point, one must note a serious trade-off between the demands that a comprehensive deliberative account seeks to impose on a decision procedure and the demand of strategy-proofness.

The strategic equivalence of the premise- and conclusion-based procedures

I have noted a curious aspect of the tenure example above. If the committee members all act on their assumed incentives in that example – some reveal truthful judgments, some strategically misrepresent their judgments, as appropriate – then the outcome under the premise- and conclusion-based procedures is exactly the same. This is strikingly at odds with the standard claim – reinforced in this article – that the premise- and conclusion-based procedures may lead to different outcomes for the same combination of individual judgments. Under what conditions does that curiosity occur?

I assumed above that the committee members are each convinced of their own judgments and have a desire to bring about collective judgments that coincide with those judgments, but that they care primarily about the final tenure decision itself. Let me introduce a general conceptual distinction (Dietrich and List 2004a). Consider again a conjunctive or disjunctive decision problem. An individual may not be able to bring about collective judgments that coincide with her own judgments on *every* proposition. If she cannot, she may care particularly about certain propositions and less about others. The individual has *outcome-oriented preferences* if her desire is first and foremost to bring about a collective judgment on the conclusion that coincides with her own individual judgment on that conclusion. This was the assumption in the tenure example. By contrast, the individual has *reason-oriented preferences* if her desire is first and foremost to bring about collective judgments on the premises that coincide with her own individual judgments on those premises. This is perhaps the motivation that deliberative democrats would like people to exhibit.

Given outcome-oriented preferences, as assumed in the tenure example, it can be proved that the premise- and conclusion-based procedures are *strategically equivalent* (Dietrich and List 2004a): Under the premise-based procedure, an individual with outcome-oriented preferences has an incentive to vote for every premise if she accepts the conclusion and against every premise otherwise. If all individuals adopt this strategy, then the premise-based procedure always generates the same decision on the conclusion as the conclusion-based one, whose outcome is the result of a truthful vote on the conclusion. Paradoxically, for outcome-oriented preferences, the mechanical differences between the two procedures are ‘cancelled out’ by their behavioural differences.

By contrast, given reason-oriented preferences, the two procedures are not strategically equivalent. Also, given reason-oriented preferences, individuals will not misrepresent their judgments under the premise-based procedure, as they seek to bring about collective judgments on the premises that coincide with their own individual judgments on those premises. It might be speculated that a public deliberative setting is more likely to induce reason-oriented preferences, while a minimal liberal setting is more likely to induce outcome-oriented preferences. If that speculation were correct, the mechanical differences between the premise- and conclusion-based procedures would not

be ‘cancelled out’ by their behavioural differences, but perhaps rather reinforced by them; furthermore, neither decision procedure would be faced with significant problems of strategic manipulation. But this thought still remains a speculation.

VIII. CONCLUDING REMARKS: TOWARDS A UNIFIED ACCOUNT

Overview

In this article, I have introduced the minimal liberal and comprehensive deliberative accounts of collective decision-making and argued that they can be associated with two different decision procedures: the conclusion- and premise-based ones. I have compared the two procedures in several respects. I have investigated when and how often they disagree, assessed them in terms of general demands on collective decision-making and shown which demands they satisfy and which not. I have also compared their truth-tracking performance in decisions on factual matters and addressed the incentives they create for decision-makers. What are the key differences between the two procedures?

The size of the public agenda. The public agenda is usually specified more narrowly for the conclusion-based procedure, while it is specified more broadly for the premise-based one.

The order of priority among the propositions. The conclusion-based procedure prioritizes the conclusion: it forms a collective judgment directly and only on the conclusion on the basis of the individual judgments on that conclusion. The premise-based procedure prioritizes the premises: it forms collective judgments directly on the premises on the basis of the individual judgments on those premises; it then derives the collective judgment on the conclusion from those collective judgments on the premises.

The strategic incentives created by the two procedures. The conclusion-based procedure is strategy-proof, while the premise-based procedure is vulnerable to the strategic misrepresentation of individual judgments.

In this final section, I revisit these key differences between the premise- and conclusion-based procedures and suggest that they determine under what conditions one procedure is more adequate than the other. In particular, I ask whether instead of associating the two procedures with two diametrically opposed accounts of decision-making – the minimal liberal and comprehensive deliberative ones – both procedures can

be embedded in a unified account in between the two ideal-typical extremes. Although I am unable to offer such an account here, I hope to sketch some of its likely elements.

The size of the public agenda

The minimal liberal and comprehensive deliberative accounts in their most extreme forms specify the public agenda in opposite ways: one makes it as narrow as possible, the other as broad as possible. But the appropriate public agenda may vary from one decision problem to another, and hence the appropriate decision procedure may also vary. What is the appropriate public agenda for a given decision problem?

For almost every decision problem, certain types of propositions, namely those on practical actions or policies, need to be included in the public agenda. But sometimes it may also be appropriate to include propositions on supporting reasons or other background considerations. As noted above, a Rawlsian criterion for the inclusion of such propositions might be whether they are ‘political’ rather than ‘metaphysical’ ones: propositions capable of being supported by an overlapping consensus within the relevant group as opposed to ones on which there is irresolvable disagreement based on different worldviews. When a wide range of people endorse certain reasons – even if they endorse them on the basis of different worldviews – then the public resolution of these reasons may enhance the legitimacy and stability of a decision. On the other hand, when reasons are controversial and divisive, their public resolution may have the opposite effect.

Should factual propositions always be included in the public agenda? If a proposition is clearly factual and its resolution poses no great difficulties, then the answer is probably positive. But whether a proposition is factual may itself be controversial. While some propositions may be unambiguously factual, others are considered factual by some people and value-laden or non-factual by others. In the global warming example, all propositions are factual ones; the first premise concerns a contingent fact about the state of the world, the second a general physical mechanism; still, some critics might describe these propositions as value-laden. In the tenure example, the propositions about the candidate’s performance in research and teaching both involve a complex combination of facts and values. In the habitat protection example, the status of the premises is unclear; there may be substantial disagreement on whether they are factual or not.

On many kinds of propositions – such as ethical or aesthetic ones – we can distinguish between *cognitivist* and *non-cognitivist* views. Cognitivists ascribe truth-values to these propositions, while non-cognitivists interpret them as expressions of attitudes or sentiments. Disagreements between cognitivists and non-cognitivists may affect not only the question of where to draw the line between factual and non-factual propositions, but also the question of what propositions to include in the public agenda. Those who hold a cognitivist view on a larger class of propositions may endorse a larger public agenda; those who hold a non-cognitivist view on larger class of propositions may endorse a smaller one.

Finally, even when people agree that a proposition is factual, they may still be so divided over that proposition's truth-value that it is better to exclude the proposition from the public agenda. Consider, for example, a proposition about the existence of God. Many theists and atheists share the belief that this proposition is factual, yet they are deeply divided over its truth-value.

I conclude that an important element of a unified account of collective decision-making will be a set of criteria that size the public agenda differently for different decision problems.

The order of priority among the propositions

As noted above, the conclusion- and premised-based procedures – in line with the minimal liberal and comprehensive deliberative accounts – prioritize the propositions differently. The order of priority among the propositions matters in at least two respects.

First, it determines whether or not the reasoning process over the propositions – the drawing of logical inferences – is 'collectivized' and, if it is, what the precise structure of that reasoning process is. Different decision problems may require a different collective reasoning process (if any) and therefore a different decision procedure.

Second, the order of priority among the propositions may affect a decision procedure's capacity to track the truth on the propositions. A well chosen order of priority – such as one that gives priority to propositions for which Condorcet's competence assumption is met by the decision-makers – can improve that capacity. The order of priority may also affect the relative likelihood of false positive and false negative

decisions and may thus be relevant to the particular asymmetry between false positives and false negatives in a given decision problem. As the ‘epistemically appropriate’ order of priority may vary from one decision problem to another, the appropriate decision procedure may also vary.

A key element of a unified account of collective decision-making will therefore be not only a set of criteria for identifying the appropriate order of priority among the propositions for a given decision problem, but also a method for making decisions based on that order. A decision procedure along these lines can be stated as follows.

A sequential priority decision procedure. First, an *order of priority* among the propositions in the public agenda is specified. Earlier propositions are interpreted as ‘prior to’ later ones, capable of serving as reasons for or against later ones. Second, the group considers the propositions in that order. For each proposition under consideration, if the proposition is unconstrained by propositions considered earlier, then the group decides the new proposition by majority voting (or some other single-proposition decision rule); but if it is constrained by propositions considered earlier (such as a conclusion constrained by earlier premises), then the group derives its judgment on the new proposition from its judgments on those earlier propositions (List 2004).

The premise- and conclusion-based procedures are both special cases of a sequential priority decision procedure, albeit for a different specification of the public agenda. The conclusion-based procedure is a sequential priority procedure applied to a simple public agenda containing only the conclusion, where that conclusion is (trivially) prioritized. The premise-based procedure is a sequential priority procedure applied to a non-simple public agenda of a conjunctive or disjunctive decision problem, where the premises and the appropriate logical relation are given priority over the conclusion.

A unified account that allows a flexible specification of both the public agenda and the order of priority among the propositions may therefore adequately cover *both* decision problems in which a premise-based approach is more adequate *and* ones in which a conclusion-based approach is more adequate.

The strategic incentives created by the two procedures

In recommending what decision procedure to use for a given decision problem, a unified account of decision-making should also take into account whether any mechanisms are in place for inducing decision-makers to reveal their judgments truthfully. These mechanisms might be incentive-based or psychological, as discussed above, and different contexts may require different such mechanisms. Clearly, more research is required here.

The non-invariance of collective decisions under changes in the public agenda and under changes in the order of priority among the propositions

Finally, let me observe an important property that a unified account of collective decision-making will exhibit: For the same underlying structure of individual judgments, the collective judgments generated by an appropriate decision procedure will not generally be invariant under changes in the public agenda and under changes in the order of priority among the propositions. To illustrate this point, consider three different decision problems with the same underlying structure of individual judgments as shown in Table 2 above.

Table 2 (restated)

	R ₁	R ₂	R ₃	C ↔ (R ₁ or R ₂ or R ₃)	C
Individual 1	True	False	False	True	True
Individual 2	False	True	False	True	True
Individual 3	False	False	True	True	True
Majority	False	False	False	True	True

The first decision problem is the habitat protection case, as discussed above. The second is the case of a three-member court that has to reach a verdict in a criminal case. The defendant might have committed one of three different crimes (R₁, R₂ or R₃). Each judge thinks that the defendant is guilty of one crime, and hence that he should go to prison (C); but the judges disagree on which crime it is. Nonetheless, the judges unanimously agree that the defendant should go to prison. The third problem is that of another three-member court that has to decide whether to repeat a trial (C) (Pettit 2001a). The trial should be repeated if and only if at least one of three procedural flaws has occurred: some of the evidence was inadmissible (R₁), the confession was forced (R₂), or some other rule of due process was violated (R₃). All judges think that the trial was procedurally flawed and that it should be repeated, although they disagree on how exactly it was flawed.

How might a unified account address these three decision problems?

The public agenda. In the habitat case, as argued above, there may be good reasons for excluding the premises from the public agenda and including in it only the conclusion. In the prison sentence case, there is every reason for including all the premises and the conclusion in the public agenda, as the court needs to provide a legal justification for its verdict. In the retrial case, again there seem to be good (though perhaps slightly weaker) reasons for including all the premises and the conclusion in the public agenda.

The order of priority among the propositions. Again, as suggested above, in the habitat case, it may be appropriate to prioritize the conclusion. In the prison sentence case, there seem to be strong reasons for prioritizing the premises, first, because they serve as the normatively significant reasons, and second, because this reduces the risk of a false positive decision – which should be reduced to protect the defendant. In the retrial case, the order of priority is less clear. But there may be good reasons for prioritizing the conclusion rather than the premises in this case, as this reduces the risk of a false negative decision – which should be reduced here to protect the defendant.

If the public agenda and the order of priority are specified as I have suggested, then a sequential priority decision procedure, applied to each case, will implement a conclusion-based approach in the habitat case, a premise-based one in the prison sentence case, and a conclusion-based one in the retrial case. So, for the same *structure* of individual judgments, the outcome is the acceptance of the conclusion in the habitat and retrial cases, and its rejection in the prison sentence case. But even in the habitat and retrial cases the conclusion is accepted through a different route. In the habitat case, it is accepted because the public agenda contains only the conclusion, which is thereby automatically prioritized. In the retrial case, it is accepted because, although the premises are also contained in the public agenda, the conclusion is prioritized so as to avoid a false negative decision.

In conclusion, the present considerations show that the structure of individual judgments underdetermines the appropriate decision at the collective level. An adequate decision procedure will neither aggregate these judgments mechanically nor systematically prioritize either the premises or the conclusion. Rather, its decision will

depend on the nature of the public agenda and on the appropriate order of priority among the propositions. In short, its decision will depend on the sphere of public reason appropriate for the decision problem in question. Often the question of what that sphere of public reason is will itself be under debate. But I hope to have shown that an answer to this question is important for a resolution of the discursive dilemma.

APPENDIX

Tables for the sketch proof of proposition 1

Table A1: An odd number of individuals $n = 3+2m$, where $m \geq 0$

	R_1	R_2	$C \leftrightarrow (R_1 \& R_2)$	C
Individual 1	True	True	True	True
Individual 2	True	False	True	False
Individual 3	False	True	True	False
Individuals 4, ..., $3+m$	True	False	True	False
Individuals $3+m+1$, ..., $3+2m$	False	True	True	False
Majority	True	True	True	False

Table A2: An even number of individuals $n = 6+2m$, where $m \geq 0$

	R_1	R_2	$C \leftrightarrow (R_1 \& R_2)$	C
Individuals 1, 2	True	True	True	True
Individuals 3, 4	True	False	True	False
Individuals 5, 6	False	True	True	False
Individuals 7, ..., $6+m$	True	False	True	False
Individuals $6+m+1$, ..., $6+2m$	False	True	True	False
Majority	True	True	True	False

Tables A1 and A2 can be modified for a public agenda containing only R_1 , R_2 , $R_1 \& R_2$ and their negations. For that agenda, delete the row corresponding to $C \leftrightarrow (R_1 \& R_2)$ from Tables A1 and A2, and replace proposition C with $R_1 \& R_2$. Tables A1 and A2 can also easily be modified for a disjunctive decision problem or a public agenda containing only R_1 , R_2 , R_1 or R_2 (or only R_1 , R_2 , $R_1 \rightarrow R_2$) and their negations.

Table A3: $n = 4$

	R_1	R_2	$C \leftrightarrow (R_1 \& R_2)$	C
Individual 1	True	True	True	True
Individual 2	True	False	True	False
Individual 3	False	True	True	False
Individual 4	True	True	False	False
Majority	True	True	True	False

Table A4 can be modified for a disjunctive decision problem, but not for a public agenda containing only R_1 , R_2 , $R_1 \& R_2$ (or $R_1 \rightarrow R_2$, or R_1 or R_2) and their negations. To obtain an impossibility result for that agenda, the integrity demand must be slightly strengthened, so as to impose a deductive closure demand. Finally, modifications of all tables for more than two premises are also straightforward.

Sketch proof of the results on the conditions under which the premise-based procedure outperforms the conclusion-based procedure

I sketch a proof for a conjunctive decision problem; the proof for a disjunctive decision problem is analogous. First consider the premise-based procedure. Each individual satisfies Condorcet's competence assumption for each premise R_i , so the group's (positive and negative) reliability on each of R_1, \dots, R_k approaches 1 with increasing group size. Therefore the probability that all of R_1, \dots, R_k are correctly decided by a majority – the product of the group's reliability levels on R_1, \dots, R_k – also approaches 1, and the probability of a correct decision on the conclusion C under the premise-based procedure approaches 1, regardless of whether C is true or false. Further, the correct decision on C is likely to be made “for the right reasons” (Bovens and Rabinowicz 2005), as all of R_1, \dots, R_k are likely to be judged correctly by a majority. Next consider the conclusion-based procedure. Each individual judges the conjunction $R_1 \& \dots \& R_k$ (and hence C) to be true only if he or she judges all of R_1, \dots, R_k to be true. If all of R_1, \dots, R_k (and C) are true, then each individual has a probability p of judging each R_i to be true, and hence a probability p^k of judging $R_1 \& \dots \& R_k$ (and C) to be true. As p is between 0.5 and $\sqrt[k]{0.5}$, it follows that p^k is *below* 0.5. By the reverse of Condorcet's jury theorem, the probability that C is judged to be true by a majority approaches 0 with increasing group size; so the positive reliability of the group decision on C under the conclusion-based procedure approaches 0. If at least one of R_1, \dots, R_k (and thus C) is false, then each individual has a probability greater than 0.5 of judging at least one of R_1, \dots, R_k (and thus C) to be false. So the probability that C is judged to be false by a majority approaches 1 with increasing group size; so the negative reliability of the group decision on C under the conclusion-based procedure approaches 1.

Sketch proof of the results on the conditions under which the conclusion-based procedure outperforms the premise-based procedure

I sketch a proof for a disjunctive decision problem; the proof for a conjunctive decision problem is analogous. First consider the conclusion-based procedure. If the disjunction R_1 or \dots or R_k (and hence C) is false, then each individual has a probability p greater than 0.5 of judging that conjunction (and C) to be false. If the disjunction R_1 or \dots or R_k (and hence C) is true, then each individual has a probability p greater than 0.5 of judging it (and C) to be true. So each individual satisfies Condorcet's competence assumption for the conclusion, and hence the group's (positive and negative) reliability on C approaches 1 with increasing group size. Next consider the premise-based procedure. If the disjunction R_1 or \dots or R_k (and hence C) is false, then each individual has a probability p greater than 0.5 of judging each of R_1, \dots, R_k to be false, and so the probability that each premise and hence the conclusion will be rejected under the premise-based procedure approaches 1 with increasing group size; thus the negative reliability of the group decision on C under the premise-based procedure approaches 1. But if the disjunction R_1 or \dots or R_k (and hence C) is true, then each individual's probability p_i of judging each R_i to be true is still below 0.5. So the probability that each premise and hence the conclusion will still be *rejected* under the premise-based procedure approaches 1 with increasing group size; thus the positive reliability of the group decision on C under the premise-based procedure approaches 0. The fact that the ‘correct’ premise R_i might receive a

larger minority than any other premise does not help. Regardless of whether C is true or false, the probability that the premise-based procedure leads to the rejection of C approaches 1 with increasing group size.

REFERENCES

- Arrow, K. J. 1951. *Social Choice and Individual Values*. New York: Wiley.
- Bovens, L., and Rabinowicz, W. 2005. Democratic Answers to Complex Questions: An Epistemic Perspective. *Synthese*, forthcoming.
- Brennan, G. 2001. Collective Coherence? *International Review of Law and Economics* 21(2): 197-211.
- Buchanan, J. M., and Tullock, G. 1962. *The Calculus of Consent*. Ann Arbor: University of Michigan Press.
- Brennan, G., and Pettit, P. 2004. *The Economy of Esteem*. Oxford: Oxford University Press.
- Chapman, B. 1998. More Easily Done than Said: Rules, Reason and Rational Social Choice. *Oxford Journal of Legal Studies* 18(2): 293-330.
- Chapman, B. 2002. Rational Aggregation. *Politics, Philosophy and Economics* 1(3): 337-354.
- Dietrich, F. 2005a. Judgment aggregation: (im)possibility theorems. *Journal of Economic Theory*, forthcoming.
- Dietrich, F. 2005b. Judgment aggregation in general logics. Working paper, University of Konstanz.
- Dietrich, F., and List, C. 2004a. Strategy-proof judgment aggregation. Working paper, LSE.
- Dietrich, F., and List, C. 2004b. A model of jury decisions where all jurors have the same evidence. *Synthese* 142: 175-202.
- Dietrich, F., and List, C. 2005. Arrow's theorem in judgment aggregation. Working paper, LSE.
- Dokow, E., and Holzman R. 2005. Aggregation of Binary Evaluations. Working paper, Technion-Israel Institute of Technology.
- Dryzek, J., and List, C. 2003. Social Choice Theory and Deliberative Democracy: A Reconciliation. *British Journal of Political Science* 33(1): 1-28.

- Dworkin, R. M. 1986. *Law's empire*. Cambridge/MA: Harvard University Press.
- Estlund, D. 1994. Opinion leaders, independence and Condorcet's jury theorem. *Theory and Decision* 36: 131-162.
- Gärdenfors, P. 2005. A representation theorem for voting with logical consequences. *Economics and Philosophy*, forthcoming.
- Gehrlein, W. 1981. The expected probability of Condorcet's paradox. *Economics Letters* 7: 33-37.
- Gehrlein, W. 1983. Condorcet's paradox. *Theory and Decision* 15: 161-197.
- Gehrlein, W., and Fishburn, P. C. 1976. The probability of the paradox of voting: A computable solution. *Journal of Economic Theory* 13: 14-25.
- Goodin, R. E., and List, C. 2006. Special Majorities Rationalized. *British Journal of Political Science*, forthcoming.
- Grofman, B., Owen, G., and Feld, S. L. 1983. Thirteen theorems in search of the truth. *Theory and Decision* 15: 261-278.
- Kornhauser, L. A. 1992. Modelling Collegial Courts. II. Legal Doctrine. *Journal of Law, Economics and Organization* 8: 441-470.
- Kornhauser, L. A., and Sager, L. G. 1986. Unpacking the Court. *Yale Law Journal* 96(1): 82-117.
- Kornhauser L. A., and Sager, L.G. 1993. The One and the Many: Adjudication in Collegial Courts. *California Law Review* 91: 1-51.
- Kornhauser, L. A., and Sager, L. G. 2004. The Many as One: Integrity and Group Choice in Paradoxical Cases. *Philosophy and Public Affairs* 32(3): 249-276.
- Ladha, K. 1992. The Condorcet jury theorem, free speech and correlated votes. *American Journal of Political Science* 36: 617-634.
- List, C. 2003. A Possibility Theorem on Aggregation over Multiple Interconnected Propositions. *Mathematical Social Sciences* 45(1): 1-13; see also a correction at <http://personal.lse.ac.uk/LIST/PDF-files/PossibilityThmCorrection.pdf>
- List, C. 2004. A Model of Path-Dependence in Decisions over Multiple Propositions. *American Political Science Review* 98(3): 495-513.
- List, C. 2005. The probability of inconsistencies in complex collective decisions. *Social Choice and Welfare* 24(1): 3-32.

- List, C., Fishkin, J., Luskin, R. C., and McLean, I. 2000. Can deliberation induce preference structuration? Paper presented at the 2000 annual meeting of the American Political Science Association.
- List, C., and Pettit, P. 2002. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy* 18(1): 89-110.
- List, C., and Pettit, P. 2004. Aggregating Sets of Judgments: Two Impossibility Results Compared. *Synthese* 140(1-2): 207-235.
- Mackie, G. 1998. All Men are Liars: Is Deliberation Meaningless? In *Deliberative Democracy*, edited by J. Elster: 97-122. New York: Cambridge University Press.
- Miller, D. 1992. Deliberative Democracy and Social Choice. *Political Studies* 40(Special Issue): 54-67.
- Nehring, K., and Puppe, C. 2005. Consistent Judgement Aggregation: A Characterization. Working paper, University of Karlsruhe.
- Norris, P. 2004. *Electoral Engineering: Voting Rules and Political Behavior*. Cambridge: Cambridge University Press.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge/MA: Harvard University Press.
- Pauly, M., and van Hees, M. 2005. Logical Constraints on Judgment Aggregation. *Journal of Philosophical Logic*, forthcoming.
- Pettit, P. 2001a. Deliberative Democracy and the Discursive Dilemma. *Philosophical Issues* 11: 268-299.
- Pettit, P. 2001b. *A Theory of Freedom*. New York: Oxford University Press.
- Pettit, P. and Smith, M. 1996. Freedom in Belief and Desire. *Journal of Philosophy* 93: 429-449.
- Pigozzi, G. 2005. Collective Decision-Making without Paradoxes: A Fusion Approach. Working paper, King's College, London.
- Rawls, J. 1993. *Political Liberalism*. New York: Columbia University Press.
- Riker, W. H. 1982. *Liberalism Against Populism*. San Francisco: W. H. Freeman.
- Sunstein, C. 1994. *Political Conflict and Legal Agreement*. Tanner Lectures on Human Values, Harvard.