

Forthcoming, *Economics and Philosophy*

Aggregating Sets of Judgments: An Impossibility Result

Christian List and Philip Pettit

revised 9 October 2001

The concern of this paper is the aggregation of sets of rationally connected judgments that the members of a group individually endorse into a corresponding, collectively endorsed set of judgments. After documenting the need for various groups to aggregate judgments, we explain how this task is challenged by the "doctrinal" or "discursive" paradox. We then show that this paradox is not just an artifact of certain specific situations, but that it actually illustrates a new impossibility theorem, according to which there exists no systematic mechanism for generally solving relevant types of aggregation problems in accordance with some undemanding conditions. This new result highlights a tension between two plausible demands: on the one hand, that a group be responsive to the judgments of individual members in forming collective judgments and, on the other, that it be rational in the judgments it collectively endorses. We consider at some length how groups can deal with this problem and evade our impossibility result, and we look at some established practices whereby they manage to do so. A formal proof of the new theorem is presented in an appendix.

Suppose that the members of a certain group each hold a rational set of judgments on some interconnected questions. And imagine that the group itself now has to form a collective, rational set of judgments on those questions. How should it go about dealing with this task? We argue that the question raised is subject to a difficulty that has recently come to attention in discussion of the doctrinal paradox in jurisprudence. And we show that there is a general impossibility theorem that that difficulty illustrates. Our paper is given to the presentation of this impossibility result and an exploration of its significance. The result naturally invites comparison with Kenneth Arrow's famous theorem (Arrow 1963; Sen 1970; Arrow 1984) and we elaborate that comparison in a companion paper (List and Pettit 2002).

The paper is in four sections. The first section documents the need for various groups to aggregate the sets of judgments maintained by members on certain issues; the second presents the discursive paradox that arises with the aggregation of sets of judgments; and then the third gives an informal statement of the impossibility result that the paradox illustrates: the formal proof is presented in an appendix. The fourth, lengthier section discusses some escape-routes from the impossibility result, asking how groups can manage – and do in practice manage – to get around the implications of the result.

1. The Task of Aggregating Judgment

The task of aggregating sets of judgments – for short, aggregating judgment – should be distinguished from two distinct tasks. First, the task of aggregating people's sets of credences in respect of certain propositions, where these are tantamount to degrees of confidence that these propositions obtain. And second, the task of aggregating people's preference orderings in respect of the alternative options or candidates available in some collective choice problem.

Judgments are modelled on acts of assent or dissent, assertion or denial, and differ from credences in not allowing of *degrees* of confidence. Under our concept of judgment, there is no such thing as judging to a certain degree that something is or is not the case; while someone may be prepared to judge only that something is probable in a certain degree, a judgment in the present sense will always be an on or off affair. It may be interesting to ask how the credences of people should be aggregated in a certain area but we here restrict our attention to the aggregation of judgment. We are interested in groups which try to form collective judgments in respect of certain issues on the basis of the judgments held and expressed by individuals in respect of those issues.

The challenge involved in the aggregation of preference orderings arises wherever a group needs to make a decision about some issue on the basis of the preferences of its members over the available alternatives: the challenge may be

large or small in scale, ranging from the electoral case of a population selecting a party or a policy to the case of a group of friends selecting a restaurant at which to have a meal. The crucial difference between aggregating preference orderings and aggregating sets of judgments is this. The problem of aggregating preference orderings arises with issues that are treated as independent, the task being to produce a collective ranking of the various alternatives in each issue. The problem with aggregating sets of judgments arises with a number of issues that are treated as interconnected rather than independent, the task being to produce a rational collective set of judgments on those different issues.¹

The task of aggregating sets of judgments arises in many different contexts. Consider the sort of group that is charged by an external authority with making certain decisions according to designated principles and particularly on the basis of multiple premises. Examples of such groups are appointment and promotions committees; committees charged with deciding who is to win a certain prize or contract; trusts that have to make judgments on the basis of a trustee's instructions; associations or the executives of associations that have to justify their actions by reference to the group's charter; corporations that have to comply with policies endorsed by their shareholders; public bodies, be they bureaucratic committees or appointed boards, that have to discharge specific briefs; and governments that are more or less bound to party programs and principles.

In each of these cases the group will have to make a set of judgments on the relevant premises and on the resulting decision. The promotions committee, for instance, will have to make a judgment on whether a candidate is formally qualified for the promotion, on whether the candidate has sufficient experience, on whether he or she is likely to establish good working relations with his or her colleagues after the promotion, and finally whether he or she should be promoted. And so on in other cases. In each case the group will naturally want to

make a set of judgments that satisfies twin conditions. On the one hand, it will want the judgments reached to be responsive to the judgments of members. On the other hand, the group will want the judgments to constitute a rational set of judgments in themselves.

For other examples of the aggregation of judgment consider purposive groups more generally, whether or not they are formally bound by certain external, antecedently designated principles. It may still be an important aspiration in these groups that members should find common principles by which to justify whatever line they collectively take on any particular issue. Think of the political party that has to work out a program for government; or the association that has to decide on the terms of its constitution; or the church that has to give an account of itself in the public forum; or the learned academy that seeks a voice in the larger world of politics and journalism. In such cases members of the group may not be presented by an external authority with antecedently designated principles on the basis of which to justify particular judgments. But the need to answer as a group to individuals in their own ranks, or to bodies in the outside world, will support a wish to reach judgments according to general principles rather than in an ad hoc fashion. For every particular issue that it faces, the group will want to be able to make a judgment on that issue that is rationally connected with other judgments that it simultaneously makes according to those principles; and it will want to be able to do this in a manner that reflects the judgments of members on those matters.²

These examples of how groups are required to aggregate judgment will suffice for purposes of introducing the problem we address in this paper. But it is worth noting that a very similar problem arises for any group that claims to promote a purpose (Pettit 2001a, ch. 4; Pettit 2001b). Such a group may only rarely have to try and resolve a set of issues that arise at one and the same time, in the manner of the examples just mentioned. But it will routinely find itself

faced with an issue on which judgments previously endorsed – and now a matter of record – have a rational bearing. Suppose that a political party announces in January, say on the basis of majority vote among its members, that it will reduce taxes if it gets into government. Suppose that it announces in June, again on the basis of majority vote, that it will increase defence spending. And now imagine that it faces the issue in September as to whether it will reduce government spending in other areas of policy or organisation. The connection between this issue and the ones previously resolved means that it will have to resolve that issue in the context of sustaining or revising its judgments on the other issues as well. It will also face a task of aggregating the sets of judgments that its members make on those different issues.

2. The Paradox in Aggregating Judgment

We now introduce a paradox that affects the aggregation of judgment. This was identified by scholars in law and economics when they drew attention to a problem that affects multi-member courts as a result of their having to decide certain cases – according to legal doctrine – on the basis of a resolution of related issues. (Kornhauser and Sager 1986; Kornhauser 1992; Kornhauser 1992; Kornhauser and Sager 1993; Chapman 1998; Chapman 1998; Brennan 2001). The problem has sometimes been described in this literature as the doctrinal paradox.

A good example of the doctrinal paradox is provided by this simple case where a three-judge court has to decide on whether a defendant is liable under a charge of breach of contract (Kornhauser and Sager 1993, 11). According to legal doctrine, the court should find against the defendant if and only if it finds, first that a valid contract was in place, and second that the defendant's behaviour was such as to breach a contract of that kind. Now imagine that the three judges, 1, 2 and 3, vote as follows on those issues and on the derivable matter of whether the

defendant is indeed liable. The 'yes' or 'no' on any row represents the disposition of the relevant judge to accept or reject the corresponding premise or conclusion.

	Valid contract?	Breach?	Liable?
1	yes	no	no
2	no	yes	no
3	yes	yes	yes

Matrix 1

There are two ways in which the court might in principle make its decision in a case like this. It might have the judges do their individual reasoning and then aggregate their votes on the conclusion – the liability issue – on, say, a majority basis. Since the conclusion does not command majority support, the defendant would go free. Call this the conclusion-driven approach. Or the court might have the judges aggregate their votes on the individual premises – the contract and breach issues – and let the resulting, collective judgments on those premises determine what it rules on the conclusion. Since each premise commands majority support, the result in this case is that the defendant would be found liable. Call this the premise-driven approach. The doctrinal paradox, as presented in the jurisprudential literature, consists in the fact that the two procedures described yield different outcomes.

This sort of paradox will arise, not just when legal doctrine dictates that certain considerations are conceptually or epistemically prior to a certain issue – an issue on which a conclusion has to be reached – and that judgments on those considerations ought to dictate the judgment on the conclusion. It arises more generally whenever a group of people discourse together with a view to forming an opinion on a certain matter that rationally connects, by the lights of all concerned, with other issues. It constitutes a discursive, not just a doctrinal, dilemma.

For an example that is close to the case just discussed, consider an issue that might arise in a workplace, among the employees of a company: say, for simplicity, a company owned by the employees (Pettit 2001b). The issue is whether to forego a pay-rise in order to spend the money thereby saved on introducing a workplace safety measure: perhaps a guard against electrocution. Let us suppose for convenience that the employees are to make the decision – perhaps because of prior resolution – on the basis of considering three separable issues: first, how serious the danger is; second, how effective the safety measure that a pay-sacrifice would buy is likely to be; and third, whether the pay-sacrifice is bearable for members individually. If an employee thinks that the danger is sufficiently serious, the safety measure sufficiently effective, and the pay-sacrifice sufficiently bearable, he or she will vote for the sacrifice; otherwise they will vote against. And so each will have to consider the three issues and then look to what should be concluded about the pay-sacrifice.

Imagine now that after appropriate dialogue and deliberation the employees are disposed to vote on the relevant premises and conclusion in the pattern illustrated by the following matrix for a group of three workers. The numerals 1, 2, and 3 represent the three employees and the ‘yes’ or ‘no’ on any row represents the disposition of the relevant employee to accept or reject the corresponding premise or conclusion.

	Serious danger?	Effective measure?	Bearable loss?	Pay-sacrifice?
1	yes	no	yes	no
2	no	yes	yes	no
3	yes	yes	no	no

Matrix 2

If this is the pattern in which the employees are inclined to vote, then a different decision will be made, depending on whether the group judgment is driven by how members judge on the premises or by how they judge on the

conclusion. Looking at the matrix, we can see that though everyone individually rejects the pay-sacrifice, a majority supports each of the premises. If we think that the views of the employees on the conclusion should determine the group-decision, then we will say that the group-conclusion should be to reject the pay-sacrifice: there are only 'no's' in the final column. But if we think that the views of the employees on the premises should determine the group-decision, then we will say that the group-conclusion should be to accept the pay-sacrifice: there is a majority of 'yes's' in each of the premise columns.

The doctrinal or discursive paradox arises in the cases considered because, while each premise commands a majority, the majorities involved are distinct and the intersection of the majorities — those who will vote for the conclusion, given they vote for all the premises — is itself a minority in the group as a whole; in this last case it has no members at all, in the previous case it has just one member, person 3. The paradox generalises in many different ways. It may arise with disjunctive as well as conjunctive reasoning, since a disjunctive set of premises, P or Q , will support a conclusion R just in case the conjunctive set of premises, $\neg P$ and $\neg Q$, supports the negation of that conclusion. It may arise whenever there are three or more propositions involved. It may arise whenever there are three or more persons involved. And so on.

The important lesson of the discursive paradox from the perspective of this paper is not the lesson emphasised in the jurisprudential literature: that there is a hard choice between letting the collective judgment on the conclusion be decided by the votes on the conclusion itself or by the votes on the premises. Rather it is the more general lesson that if we go the normal, majoritarian way in aggregating a collective set of judgments from individual sets of judgments, even from individual sets that are themselves entirely rational, then we may end up with a collective set that is irrational. While every individual in each of our examples is perfectly rational in their pattern of assent to premises and

conclusion, the collectivity would not be rational in its pattern of assent if it followed the procedure of majority voting on each of the premises *and* also on the conclusion. It would not be rational in that it would reject the conclusion in each case, while endorsing premises that support it.

A group in the sort of predicament illustrated faces a hard choice, as the jurisprudential literature says, but the options are more general than those associated with privileging the conclusion-votes and privileging the premise-votes. The group may allow the votes of the members on each of a rationally connected set of issues to determine the view of the collectivity on that issue, in which case there is a risk of the collectivity's holding an irrational set of views. Or the group may take steps to ensure that the collective view espoused will be rational, in which case it may be necessary to ignore the vote of the majority on some issue. We may describe the first alternative as that of ensuring individual responsiveness – that is, the responsiveness of the collective view to individual votes – and the second as that of ensuring collective rationality.

Enfranchising the votes of individuals on the conclusion, where that is part of a broader project of enfranchising the votes of individuals on each of the issues raised in the premises and the conclusion, means ensuring individual responsiveness. Enfranchising the votes of individuals on the premises, and letting the collective views thereby determined fix the collective conclusion, is an instance of ensuring collective rationality. But while it is an instance of that strategy, it is not the only way of pursuing it. The members of the group might ensure collective rationality, not by rejecting the view taken by majority vote on the conclusion, but rather by rejecting the view taken by majority vote on one of the premises. They might be open to the possibility of practising *modus tollens* instead of *modus ponens*.

The discursive paradox brings out the tension between two plausible demands that we might want to make on the aggregation of judgment. The first

is the demand that in aggregating judgment a group should be responsive to the views of members on each of the judgments involved. The second is the demand that in aggregating judgment a group should reach a collective set of judgments that is itself rational. The paradox shows that the two demands are sometimes in conflict, so that a group that tries to aggregate judgment faces a dilemma. It may seek to be responsive to the judgments of individuals, thereby risking collective rationality. Or it may seek to ensure collective rationality, thereby risking responsiveness to the judgments of individuals. The discursive paradox is a discursive dilemma.

3. A Simple Impossibility Result on the Aggregation of Judgment

The cases of the discursive paradox that we considered show that if we take the sets of views held among a group of people on a range of propositions then, even if those sets of views each satisfy rationality constraints — even if they are each consistent and the like — the set of collective views derived by majority voting on each issue may not satisfy such rationality constraints. This suggests that a procedure like systematic majority voting — majority voting on each issue — can never be guaranteed against producing an irrational set of collective judgments.

These intuitive considerations point us towards the prospect of a general proof that any procedure that is akin to majority voting in certain ways — that satisfies various conditions, yet to be identified, that are exemplified in majority voting — will be vulnerable to this paradoxical possibility. They suggest the prospect of establishing the impossibility of finding a procedure for aggregating judgment that satisfies such conditions *and* that is proof against paradox. We present the intuitive elements of our result in this section; a formal proof is given in an appendix.

Let us suppose that we have a set of individuals, N , each named by a numeral: $1, 2, 3, \dots, n$; we assume that N has at least two members. And let X be

the set of propositions on which judgments are to be made, including atomic propositions like P and Q and compound propositions such as $(P \wedge Q)$, $(P \vee Q)$, $((P \wedge Q) \rightarrow \neg R)$, and so on; we assume that X contains at least two distinct atomic propositions, P , Q , their conjunction, $(P \wedge Q)$, and the negation of their conjunction $\neg(P \wedge Q)$.

We can identify the set of judgments made by an individual, i , in respect of X – we call them i 's *personal set of judgments* – with the set of those propositions in X to which he or she assents. Notice that this means that dissent – saying ‘no’ or ‘false’ to a proposition – does not involve making a judgment in respect of it; person i will make a judgment, strictly speaking, only if he or she also says ‘yes’ or ‘true’ to the negation of the proposition. The individual's personal set of judgments will thus be a subset Φ_i of X . With this notion secure, we can define the *profile of personal sets of judgments* that are held across the group, N , as the n -tuple of those personal sets of judgments: $\{\Phi_i\}_{i \in N}$.

A personal set of judgments Φ_i may or may not satisfy rationality constraints. We assume that all personal sets of judgments in the group, N , satisfy three such constraints. We describe these, in standard terminology, as completeness, consistency and deductive closure. A personal set of judgments Φ_i will be *complete* if, for all propositions ϕ in X , either ϕ or its negation $\neg\phi$ (or unnegated form ψ in case $\phi = \neg\psi$) is contained in the set Φ_i to which the person assents. It will be *consistent* if there is no proposition ϕ in X such that both ϕ and $\neg\phi$ are contained in Φ_i .³ And it will be *deductively closed* if, whenever Φ_i logically entails some other proposition φ in X , φ is also contained in Φ_i .

With these definitions and assumptions in place, the question that arises is whether there is any procedure of aggregation such that from the complete, consistent and deductively closed personal sets of judgments that individuals endorse it will allow the group to derive a collective set of judgments that is also complete, consistent and deductively closed. In particular, the question is

whether there is any procedure for doing this that satisfies some minimal conditions.

The lesson of the discursive paradox is that a majoritarian procedure will fail under certain profiles of personal sets of judgments to yield a collective set of judgments that is deductively closed and consistent. Consider the following table, which is closely related to the first matrix used above in illustrating the discursive paradox. It tells us what each person's judgment is in respect of each of the propositions in the top row and also what the majority judgment is on each of those propositions.

	P	Q	$(P \wedge Q)$	$\neg(P \wedge Q)$
Person 1	true	true	true	false
Person 2	true	false	false	true
Person 3	false	true	false	true
Majority	true	true	false	true

Table 1

In the scenario described in this table, each personal set of judgments satisfies completeness, consistency and deductive closure but the group will fail to do so under a majoritarian procedure of aggregation. A majority supports P and a majority supports Q , but $(P \wedge Q)$ does not get majority support. That means that the collective set of judgments is not deductively closed; it does not include the judgment that $(P \wedge Q)$, though that proposition is entailed by other propositions in the set: that is, P and Q . And not only that. The proposition $\neg(P \wedge Q)$ is in the group's collective set of judgments. So if the group's procedure is to endorse the *deductive closure* of its majority judgments (rather than just those majority judgments) as its collective set of judgments, then the resulting

collective set of judgments will be inconsistent: it contains P , Q , $(P \wedge Q)$ and $\neg(P \wedge Q)$.

So a majoritarian procedure for aggregating judgment will fail under some profiles of personal sets of judgments to generate a collective set of judgments that score like the personal sets in the properties of completeness, consistency and deductive closure. But is the problem tied to the use of the majoritarian procedure? Or can we identify features of majoritarianism such that *any* aggregation procedure with those features will fail in the same way? Our theorem shows that the answer to this question is positive.

Take a procedure or function F , to be called a *judgment aggregation function*, whose input is a profile of personal sets of judgments, $\{\Phi_i\}_{i \in N}$, and whose output is a corresponding *collective set of judgments*, Φ , to be endorsed by the group as a whole. The argument is that no F that satisfies certain minimal conditions — conditions that majoritarianism and many other aggregation procedures satisfy — can generate collective judgments that display completeness, consistency and deductive closure.

The minimal conditions that we specify are these.

UNIVERSAL DOMAIN (U). A judgment aggregation function, F , should accept as admissible input any logically possible profile of personal sets of judgments — any logically possible n -tuple $\{\Phi_i\}_{i \in N}$ — provided each person's set of judgments, Φ_i , satisfies the conditions of completeness, consistency and deductive closure.

ANONYMITY (A). The collective set of judgments, Φ , that is yielded by F should be invariant under any permutation of the individuals in N . (Intuitively, this requires that people should be treated in an even-handed way by the aggregation function; no one individual's judgments should be given special weight in determining the collective judgments.)

SYSTEMATICITY (S). For any two propositions ϕ and ψ in X , if every individual in N makes exactly the same judgment (acceptance/rejection) on ϕ as he or she

makes on ψ , then the collective judgment (acceptance/rejection) on ϕ should also be the same as that on ψ , and the same pattern of dependence of collective judgments on individual ones should hold for all profiles in the domain of F . (Intuitively, this requires that propositions or issues should be treated in an even-handed way by the aggregation function; the collective judgment on each proposition should depend exclusively on the pattern of individual judgments on that proposition. In particular, the collective judgment on no proposition should be given special weight in determining the collective judgments on others.)

The condition of universal domain is scarcely contestable. An aggregation procedure could hardly be described as fully satisfactory if its success depended on the particular profile of personal sets of judgments that happens to be presented for aggregation. If it is to be fully satisfactory, then it must work for any logically possible profile of personal sets of judgments. So at any rate it seems natural to think.

The condition of anonymity is also compelling. Anonymity stipulates that no one's identity ought to matter in determining the result that the aggregation procedure yields. It may be important, for example, that a majority of such and such a size supports a given proposition but it cannot be important that a majority with such and such members does so.

The condition of systematicity says, somewhat more controversially, that the collective judgment on any proposition should depend only on the personal judgments in respect of that proposition and that there should be no difference between the ways in which different collective judgments depend on relevant personal judgments.⁴ We can express this by saying that if two propositions have support from just the same people, then they ought to attract the same collective response. Combined with anonymity, that implies that if two propositions have the same degree of support, though perhaps from different people, then they

ought to attract the same collective response: if majoritarianism rules in one case, it ought to rule in the other, and so on.

These conditions are minimal in the sense that they will be satisfied by a whole range of imaginable aggregation procedures. Not just by majoritarianism, intuitively characterised, but also by unanimitarianism and minoritarianism and any procedures requiring special majorities of two-thirds and the like. But minimal though they are, the formal proof in appendix 1 establishes the following theorem:⁵

Theorem 1. There exists no judgment aggregation function F generating complete, consistent and deductively closed collective sets of judgments which satisfies universal domain, anonymity and systematicity.

One last comment. We constrain the profiles on the input side of a judgment aggregation function by stipulating that they contain only personal sets of judgments that are complete, consistent and deductively closed. This is a tough constraint that will not often be met in practice. But our theorem is strengthened, not weakened, by the stipulation. If the aggregation even of such well-behaved sets of personal judgments causes trouble then, *a fortiori*, the aggregation of less well-behaved sets may be expected to do so; logically, our theorem will still hold if the domain defined in condition (U) is extended.

4. Strategies for Evading the Impossibility

If our argument in the first section is sound, then groups are routinely required to aggregate judgment and so there is practical interest in the question of how they manage to avoid the problem identified in our impossibility result. If groups are going to be a significant presence in social life – in particular, if they are to be able to be rationally devoted to a range of different purposes – then they must succeed in making and upholding collective sets of judgments that satisfy constraints like consistency and deductive closure. So how do they manage to do this, given that any procedure that satisfies our conditions of

universal domain, anonymity and systematicity is bound to fail them in some cases?

The answer is that, strictly speaking, they don't manage to do this. They ensconce the rule of rationality in their collective lives, as is required of them, only at a cost in the extent to which they are responsive to the views of their members: in short, only at a cost to universal domain, anonymity or systematicity. Or they ensconce not quite the rule of rationality – not quite the rule of completeness, consistency and deductive closure – but only something that roughly approximates it. There are six general types of strategy for 'collectivising reason', as we might put it, and they correspond, on the one hand, to breaches of the three conditions of universal domain, anonymity and systematicity; and, on the other, to shortfalls in completeness, consistency and deductive closure. Of these strategies, the first four are capable of being operationalised, the last two are not really feasible. We review them briefly in the remainder of this section. (For a comparison between these strategies and strategies for avoiding the Arrowian impossibility theorem, see List and Pettit 2002).

Relaxing universal domain: the convergence strategy

We mentioned in discussion of the discursive paradox that the problem identified there arises because while each premise attracts majority support, those majorities diverge widely. The majorities that support the premises in our examples of the paradox are such that the intersection between them is a minority in the group as a whole. If the intersection of the majorities on all premises were itself a majority – in which case, of course, the intersection of the majorities on all propositions, premises and conclusion, would be a majority – then the resulting profiles of personal sets of judgments would not give rise to the discursive paradox.

One suggestion might be that a group can ensure that this condition is met by encouraging convergence, whether on the basis of increased interpersonal deliberation – if that is thought effective in building convergence (but see Sunstein 2000) – or by some other means. If a group can moderate the divergence in views between its members such that the intersection of the majorities on all propositions in any relevant set becomes itself a majority, then the paradox will be avoided. With regard to this highly restricted set of profiles – as distinct from the unrestricted set allowed under universal domain – majoritarian voting will constitute a judgment aggregation function that generates a complete, consistent and deductively closed collective set of judgments and one that satisfies anonymity and systematicity. (Compare Miller 1992; Dryzek and List 1999; and List, McLean, Fishkin and Luskin 2000).

This suggestion raises the question as to whether we can identify any structural condition on the profiles of individual judgments in respect of a set of propositions such that if individuals converge in satisfying that condition, then the impossibility problem will not occur. Should we be able to identify the sort of condition envisaged then we could be sure that, starting from any profile of individual judgments satisfying the condition, a straightforward procedure like majoritarian voting on each proposition would yield a collective set of judgments that is consistent, complete and deductively closed.

As it happens, there is an identifiable condition that plays this role (List 2001): unidimensional alignment. A profile of sets of individual judgments is unidimensionally aligned if the individuals involved can be linearly ordered from left to right, as in Table 3, so that for each proposition to be voted on, the individuals holding it to be true lie all to the right, or all to the left, of those holding it to be false. The impossibility result can be circumvented if the profiles admitted in the domain are all unidimensionally aligned in this sense. The

theorem, formulated so as to bring out the relationship with the impossibility theorem, is this.

Theorem 2. (List 2001) Suppose that the number of individuals is odd. On the domain of unidimensionally aligned profiles of personal sets of judgments, there exists a judgment aggregation function F generating complete, consistent and deductively closed collective sets of judgments which satisfies anonymity and systematicity.

	Person 3	Person 1	Person 2	Person 5	Person 4
P	true	true	true	false	false
Q	false	false	false	true	true
$(P \wedge Q)$	false	false	false	false	false
$\neg(P \wedge Q)$	true	true	true	true	true

Table 3

The theorem is quite intuitive. Applying propositionwise majoritarian voting to the profile in table 3 would result in a collective set of judgments – P , $\neg Q$ and $\neg(P \wedge Q)$ – that is consistent, complete and deductively closed. This is precisely the set of judgments held by Person 2, the median person with respect to the left-right alignment. So long as the number of individuals is odd, the median person will share the majority view on each proposition, and consequently the median person's personal set of judgments will be endorsed under a procedure of majority voting on each proposition. The completeness, consistency and deductive closure of the resulting collective set of judgments is a consequence of the completeness, consistency and deductive closure of the median person's set of judgments.

Relaxing anonymity: the authority strategy

One of the presuppositions in the discursive paradox is, roughly, that each person has an equal say; if any one of them had the authority to decide the group's judgments, then there would be no problem. This presupposition is registered in the second condition on our impossibility result: that of anonymity. As the relaxation of universal domain can provide an escape from the paradox and the impossibility result, so the relaxation of anonymity can provide one too. Suppose that a particular person is appointed as an authority in the sense that no

matter what the profile of personal sets of judgments, this person's set of judgments is taken as the overall collective set of judgments. This procedure constitutes a judgment aggregation function generating complete, consistent and deductively closed collective sets of judgments that satisfies universal domain and systematicity but not, of course, anonymity.

The relaxation involved here goes with the group strategy of naming one individual, or perhaps a small set of individuals, as dictator or plenipotentiary. More realistically, the plenipotentiary might have the somewhat circumscribed power of deciding how the group should go in the event of a paradox arising in the aggregation of judgments (formally, this circumscription would involve a breach of systematicity too). This power might be not just circumscribed but provisional. It might be exercised under the proviso that if a certain number of members think that it is badly exercised then they have a right to appeal – and a right to be heard by an authoritative judiciary.

As the first strategy can be associated with the convergence-building approach favoured by some deliberative democrats, this second strategy can be identified with an approach that prevails implicitly in many corporate bodies. The directors are implicitly given authority in such bodies to ensure that, however the expressed or presumed views of members go, the body as a whole should not fall into collective unreason. That authority may be more or less absolute but in the general run of things it will be subject to appeal and reversal; even in the commercial corporation, after all, there will always be the possibility of a shareholder revolt.

Relaxing systematicity: the priority strategy

Another condition on which our impossibility result is based is that if majority rule or some other procedure of aggregation is used to decide the collective view on some propositions, then that procedure should be used to decide the collective view on all. Why should the collective view on some

propositions be determined differently from how it is determined on others? Yet if the collective view is determined on the same basis for each proposition in a given set, that is, if we impose systematicity along with anonymity and universal domain, then we know that the collective view will often fail to satisfy rationality constraints.

The obvious response to this observation is to say that what a group ought to do, where possible, is to decide with any set of propositions – any potentially troubling set of propositions – that some of those propositions should be given priority over others and that the decision on the others should be determined by the overall decision on the prioritised propositions, not by the same procedure as in the prioritised cases. If the pattern of majority support among the prioritised propositions rationally requires a certain pattern of support among the other, 'second-grade' propositions, then that is the pattern that ought to prevail there, regardless of whether each of these propositions is accepted or rejected by a majority. In particular, suppose a group identifies a subset Y of the set of propositions X such that the propositions in Y are considered prior to the propositions outside Y and the propositions in Y are logically independent (in the technical sense that any combination of truth-values can be consistently assigned to them). If the group exercises majority voting only on the propositions in Y (jointly with some tie-breaking rule) and then endorses the *deductive closure* of these majority judgments as its collective set of judgments, then this procedure constitutes a judgment aggregation function generating consistent and deductively closed collective sets of judgments that satisfies universal domain and anonymity (but not systematicity). (Note that this procedure may not always guarantee completeness, since the majority pattern of acceptance/rejection on the epistemically prior propositions in X may not logically determine a unique pattern of acceptance/rejection on *all* other propositions; *some* propositions outside Y -- particularly if they are logically independent from the propositions

in Y -- may not be decided by the given acceptance/rejection pattern on the prioritised propositions in Y .)

Where the earlier strategies relax universal domain and anonymity, this strategy would relax systematicity. It would argue for taking a different line on how the collective view on some propositions should be determined from the line that is taken on how the collective view on others should be fixed.

The priority strategy is nicely exemplified by the premise-driven approach described earlier. This consists precisely in taking certain propositions as prior, treating them as potential premises in a *modus ponens* pattern of reasoning, and letting the collective, majority-determined views on those premises dictate the collective view to be taken on the conclusion. The priority strategy promises to be a successful way for a group to ensconce rationality at the collective level, whenever there is a prospect of agreement on which propositions to prioritise.

But such agreement, of course, will not always be forthcoming. There is no independent mark of priority attaching to those propositions that ought to be treated in the manner of premises, as distinct from those that ought to be treated as conclusions. True, the propositions that are to be treated as premises will have to be logically independent of one another – they will have to allow of any distribution of truth-values – since otherwise they will be liable to give rise in themselves to discursive paradox. But consistently with that condition, there may be a variety of different possible 'premise' sets.

What determines whether someone will see a proper subset of the propositions on which they and their group have to judge as fit to be treated as premises? People often differ in the background assumptions they make as to which sorts of propositions matter most, and for this reason they will often differ in which propositions, if any, they see as fit for treatment as premises. For one person certain atomic propositions like P , Q and R may seem most natural to be treated as premises in relation to a compound proposition, say $\neg(P \wedge Q \wedge R)$, on

which they also have to judge. For another it may seem that that compound proposition lends itself more readily to resolution than any of the atomic propositions, so that assent to the atomic propositions ought to be shaped by whether the compound proposition commands assent, not the other way around. One individual's conclusion may be another's premise; one individual's *modus ponens* may be another's *modus tollens*.

Relaxing completeness: the special-support strategy

A fourth strategy that a group might employ to get around our impossibility result would be to give up explicitly on completeness. The strategy might involve treating only those propositions that command some special degree of assent together with their deductive closure as collectively endorsed. Take for example the strategy requiring unanimous assent on any proposition that is to be collectively endorsed. This procedure would constitute a judgment aggregation function generating consistent and deductively closed sets of collective judgments that satisfies universal domain, anonymity and systematicity. It would not, of course, guarantee completeness. There might be many propositions in the set to be decided such that neither they nor their negations are collectively endorsed.

This strategy has its attractions as an escape route from our result. Let the incompleteness that it would introduce not inhibit a group from getting on with its business, whatever that is, and the strategy is bound to prove fetching. There are not many cases where the incompleteness is likely to be innocuous, however. In most situations a group that could only decide on judgment and action when members were in significant agreement about the issues involved would be unable to get anything done; it would be in a state of suspended animation.

The strategy envisaged raises the question as to how far it can work with degrees of special support that are less than unanimity and, as it happens, there is a general result that bears on the matter. Let a supermajority of more than (k -

$1)/k$ of the individuals be required, rather than anything as simple as unanimity, where k is the number of propositions to be voted on (counting each proposition/negation pair as one proposition). In other words, let the procedure whereby collective judgments are to be formed be this: the individuals vote on each proposition and a proposition is collectively endorsed only if it is supported by a supermajority. The procedure of voting by this supermajority will satisfy universal domain, anonymity and systematicity and generate consistent collective sets of judgments. While the collective sets of judgments may not satisfy completeness, deductive closure can be assured so long as systematicity is relaxed: this is unsurprising, since a proposition might be included in the relevant deductive closure without itself commanding the requisite supermajority assent. The relevant theorem can be formulated as follows.

Theorem 3. (List 2001) Suppose there are k propositions to be voted on (counting each proposition/negation pair as one proposition). Proposition-wise supermajority voting with a threshold of $(k-1)/k$ constitutes a judgment aggregation function F generating consistent (but not necessarily complete and deductively closed) collective sets of judgments which satisfies universal domain, anonymity and systematicity. Moreover, the function G , defined as the deductive closure of F , constitutes a judgment aggregation function generating consistent and deductively closed (but not necessarily complete) collective sets of judgments which satisfies universal domain and anonymity (but not necessarily systematicity).

Relaxing consistency and relaxing deductive closure: two unattractive strategies

And so, finally, to the last two types of strategy that a group might envisage. The first would take the group to be committed to all the propositions that command majority assent – or that pass whatever criterion of aggregation is employed – and to the deductive closure of those positively endorsed propositions. This would have the result that, while complete, the set of

propositions collectively endorsed might be inconsistent: they might be like the inconsistent propositions endorsed in the configuration of table 2.

The second unattractive strategy would take the group to be committed to all the propositions that command majority assent (jointly with some tie-breaking rule) – or that pass whatever criterion of aggregation is employed – but not to their deductive closure. The set of collective judgments, under this approach, would be complete and in a weak technical sense consistent; it would not contain a proposition and its negation. But, not being deductively closed, the judgments in the set might be such that were the group to endorse their implications, then it would immediately lapse into such inconsistency.

Neither of these strategies is likely to be attractive for any group that faces the need to aggregate judgment. Where the first three strategies seek to moderate the domain of individual judgments considered admissible or the extent to which the collective judgments are responsive to individual judgments, the last two strategies belong with the fourth in seeking to establish something less than the rule of rationality in the collective judgments of a group. But what the last two seek to establish, unlike the fourth, is something that departs too far from rationality to have any attraction. The fourth strategy would merely make the collective judgments of a group incomplete in the sense defined. The last two strategies would make those judgments so irrational that it is hard to see how they could serve to guide action.

Conclusion

The types of strategy just reviewed are very broad in character and allow both of variation and mixture. An exploration of the strategies at a greater level of detail, however, is beyond the scope of the present paper, as is the development of technical results that would build on the simple theorem presented here. We hope only to have indicated that our impossibility result and the conditions associated with it provide useful guidance in approaching the

question of how collectivities can manage to be relatively responsive to the judgments of their members and yet relatively rational in the judgments they collectively endorse. It turns out that while the impossibility result prevails, still there are familiar tactics that groups have established for circumventing its practical implications.⁶

Nuffield College
Oxford OX1 1NF, U.K.

christian.list@nuffield.oxford.ac.uk

Research School of Social Sciences
Australian National University
Canberra, ACT 0200, Australia
pnp@coombs.anu.edu.au

Appendix. A Simple Impossibility Theorem on the Aggregation of Judgment*

Let $N = \{1, 2, \dots, n\}$ be a set of individuals (we assume that $n \geq 2$). Let $X = \{\phi_1, \phi_2, \phi_3, \dots\}$ be a set of well-formed formulae of the propositional calculus, interpreted as the set of propositions on which simultaneous judgments are to be formed. We assume that X contains at least two distinct atomic propositions, P and Q , and their conjunction, $(P \wedge Q)$.⁷ Moreover, we assume that X contains proposition-negation pairs, specifically, whenever $\phi \in X$, then $\neg \phi \in X$ (the negation of ϕ is contained in X) or $\psi \in X$ where $\phi = \neg \psi$ (ϕ is the negation of another proposition contained in X).

To each individual i in N , there corresponds a *personal set of judgments* $\Phi_i \subseteq X$, where Φ_i is interpreted as the set of all those propositions in X to which individual i assents. Note that, if a proposition ϕ is not contained in Φ_i , this means only that it is not the case that individual i assents to ϕ ; it does not by itself mean that individual i assents to the negation of ϕ . Rather, individual i assents to the negation of ϕ only if $\neg \phi$ is contained in Φ_i (or if ψ is contained in Φ_i where $\phi = \neg \psi$).

A *profile of personal sets of judgments* is an n -tuple $\{\Phi_i\}_{i \in N}$, containing one personal set of judgments Φ_i for each individual i in N . A *judgment aggregation function* is a function F whose input is a profile of personal sets of judgments and whose

* This appendix is the work of Christian List.

output is a *collective set of judgments* $\Phi \subseteq X$, where Φ is interpreted as the set of all those propositions in X to which the group N assents.

A (personal or collective) set of judgments Φ is said to be *complete* if, for all ϕ in X , at least one of $\phi \in \Phi$, $\neg \phi \in \Phi$, or $\psi \in \Phi$, where $\phi = \neg \psi$, holds. Completeness is the requirement that, for every proposition-negation pair, *at least* one of the two be accepted. Φ is said to be *consistent* if, for all ϕ in X , at most one of $\phi \in \Phi$ or $\neg \phi \in \Phi$ holds. Consistency is the requirement that, for every proposition-negation pair, *at most* one of the two be accepted. Φ is said to be *deductively closed* if, for any φ in X , if Φ logically entails φ , then $\varphi \in \Phi$. Deductive closure is the requirement that a proposition be accepted if it is a logical implication of other accepted propositions.

UNIVERSAL DOMAIN (U). The domain of F is the set of all logically possible profiles of complete, consistent and deductively closed personal sets of judgments.

ANONYMITY (A). For any $\{\Phi_i\}_{i \in N}$ in the domain of F and any permutation $\sigma: N \rightarrow N$, $F(\{\Phi_i\}_{i \in N}) = F(\{\Phi_{\sigma(i)}\}_{i \in N})$.

SYSTEMATICITY (S). There exists a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ such that, for any $\{\Phi_i\}_{i \in N}$ in the domain of F , $F(\{\Phi_i\}_{i \in N}) = \{\phi \in X : f(\delta_1(\phi), \delta_2(\phi), \dots, \delta_n(\phi)) = 1\}$, where, for each i in N and each ϕ in X , $\delta_i(\phi) = 1$ if $\phi \in \Phi_i$ and $\delta_i(\phi) = 0$ if $\phi \notin \Phi_i$.

Theorem 1. There exists no judgment aggregation function F generating complete, consistent and deductively closed collective sets of judgments which satisfies (U), (A) and (S).

Proof.⁸ Suppose F satisfies (U), (A) and (S) and generates complete, consistent and deductively closed collective sets of judgments. By assumption, we may assume that $P, Q, (P \wedge Q), \neg(P \wedge Q) \in X$.

Step (1). Since F satisfies (S), there exists a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ such that, for any $\{\Phi_i\}_{i \in N}$ in the domain of F , $F(\{\Phi_i\}_{i \in N}) = \{\phi \in X : f(\delta_1(\phi), \delta_2(\phi), \dots, \delta_n(\phi)) = 1\}$, where, for each i in N and each ϕ in X , $\delta_i(\phi) = 1$ if $\phi \in \Phi_i$ and $\delta_i(\phi) = 0$ if $\phi \notin \Phi_i$. Since

F satisfies (A), for any $(d_1, d_2, \dots, d_n) \in \{0, 1\}^n$ and any permutation $\sigma: N \rightarrow N$, $f(d_1, d_2, \dots, d_n) = f(d_{\sigma(1)}, d_{\sigma(2)}, \dots, d_{\sigma(n)})$, and thus for any $(d_1, d_2, \dots, d_n), (e_1, e_2, \dots, e_n) \in \{0, 1\}^n$, $f(d_1, d_2, \dots, d_n) = f(e_1, e_2, \dots, e_n)$ if $|\{i \in N : d_i = 1\}| = |\{i \in N : e_i = 1\}|$. For each $\phi \in X$, define $N_\phi := \{i \in N : \phi \in \Phi_i\}$. Then, for any $\phi, \psi \in X$, if $|N_\phi| = |N_\psi|$, then $\phi \in F(\{\Phi_i\}_{i \in N})$ if and only if $\psi \in F(\{\Phi_i\}_{i \in N})$.

Step (2). By assumption, $n \geq 2$. Consider a profile of personal sets of judgments $\{\Phi_i\}_{i \in N}$ with the following properties:

	$\delta_i(P)$	$\delta_i(Q)$	$\delta_i((P \wedge Q))$	$\delta_i(\neg(P \wedge Q))$
$i = 1$	1	1	1	0
$i = 2$	1	0	0	1
$i = 3$	0	1	0	1
$i > 3$ and i is even	1	1	1	0
$i > 3$ and i is odd	0	0	0	1

Table 3

Let $\Phi = F(\{\Phi_i\}_{i \in N})$.

Case (i). n is even. We have $|N_{(P \wedge Q)}| = |N_{\neg(P \wedge Q)}|$, whence $(P \wedge Q) \in \Phi$ if and only if $\neg(P \wedge Q) \in \Phi$. By the completeness of Φ , at least one of $(P \wedge Q) \in \Phi$ or $\neg(P \wedge Q) \in \Phi$ must hold. But then we must have both $(P \wedge Q) \in \Phi$ and $\neg(P \wedge Q) \in \Phi$, which contradicts the assumption that Φ is consistent.

Case (ii). n is odd. We have $|N_P| = |N_Q| = |N_{\neg(P \wedge Q)}|$, whence either all, or none, of $P \in \Phi$, $Q \in \Phi$ and $\neg(P \wedge Q) \in \Phi$ must hold.

If $P \in \Phi$, $Q \in \Phi$ and $\neg(P \wedge Q) \in \Phi$, then, since $P, Q \in \Phi$, Φ logically entails $(P \wedge Q)$, and hence, by the deductive closure of Φ , $(P \wedge Q) \in \Phi$. But then both $\neg(P \wedge Q) \in \Phi$ and $(P \wedge Q) \in \Phi$, and Φ violates consistency, which contradicts the assumption that Φ is consistent.

If $P \notin \Phi$, $Q \notin \Phi$ and $\neg(P \wedge Q) \notin \Phi$, then, by the completeness of Φ , $(P \wedge Q) \in \Phi$, but since $(P \wedge Q)$ logically entails P , and Φ is deductively closed, we must have $P \in \Phi$, a contradiction. **Q.E.D.**

References

- Arrow, K. (1963). Social Choice and Individual Values. New York, Wiley.
- Arrow, K. (1984). Social Choice and Justice. Oxford, Blackwell.
- Bohman, J. and W. Rehg, Eds. (1997). Deliberative Democracy: Essays on Reason and Politics. Cambridge, Mass., MIT Press.
- Brennan, G. (2001). "Collective Coherence?" International Review of Law and Economics 21: 197-211.
- Chapman, B. (1998). "Law, Incommensurability, and Conceptually Sequenced Argument." University of Pennsylvania Law Review 146: 1487-1582.
- Chapman, B. (1998). "More Easily Done than Said: Rules, Reason and Rational Social Choice." Oxford Journal of Legal Studies 18: 293-329.
- Dryzek, J. and C. List (1999). "Social Choice Theory and Deliberative Democracy: A Reconciliation." European Consortium for Political Research, Mannheim.
- Elster, J., Ed. (1998). Deliberative Democracy. Cambridge, Cambridge University Press.
- Kornhauser, L. A. (1992). "Modelling Collegial Courts. I. Path-Dependence." International Review of Law and Economics 12: 169-85.
- Kornhauser, L. A. (1992). "Modelling Collegial Courts. II. Legal Doctrine." Journal of Law, Economics and Organization 8: 441-70.
- Kornhauser, L. A. and L. G. Sager (1986). "Unpacking the Court." Yale Law Journal 82.
- Kornhauser, L. A. and L. G. Sager (1993). "The One and the Many: Adjudication in Collegial Courts." California Law Review 81: 1-59.

- Miller, D. (1992). "Deliberative Democracy and Social Choice." Political Studies 40 (special issue): 54-67.
- List, C. (2001). Mission Impossible? The Problem of Democratic Aggregation in the Face of Arrow's Theorem. DPhil-thesis in Politics. Oxford, Oxford University.
- List, C., I. McLean, J. Fishkin and R. Luskin (2000). "Can Deliberation Induce Greater Preference Structuration?" Conference on Deliberating About Deliberative Democracy, University of Texas, Austin.
- List, C. and P. Pettit (2002). "Aggregating Sets of Judgments: Two Impossibility Results Compared." Synthese, forthcoming.
- Pettit, P. (2001a). A Theory of Freedom: From the Psychology to the Politics of Agency. Cambridge, Polity, and New York, Oxford University Press.
- Pettit, P. (2001b). "Deliberative Democracy and the Discursive Dilemma." Philosophical Issues 11.
- Sen, A. (1970). Collective Choice and Social Welfare. Edinburgh, Oliver and Boyd.
- Sunstein, C. (2000). The Law of Group Polarization. University of Chicago, Chicago.

¹ For the moment we say little in detail on what precisely 'rationality' requires, whether in an individual or in a collectivity; we remedy that defect in section 3, when we turn to the impossibility result proper.

² Deliberative democrats often focus on groups of this kind, urging that particular decisions should be made on the basis of commonly agreed commitments – say, commonly accepted judgments as to the shared interests of members of the group. Deliberative democrats argue that such groups ought to be inclusive, judgmental and dialogical (see Bohman and Rehg 1997; Elster 1998; Pettit 2001b). All members should be equally entitled to participate in the

decision on how to resolve relevant collective issues, or bundles of issues. Before making the actual decision, members should deliberate on the basis of presumptively common concerns about which judgments they should make. And members should conduct this deliberation in open and unforced dialogue with one another, whether in a centralised forum or in various decentralised contexts. When deliberative democrats argue for this sort of procedure, then they are arguing precisely that the groups in question should aggregate sets of judgments, not just aggregate their particular preferences.

³ Note that this is a rather weak syntactic notion of consistency. It requires only that no proposition and its negation be simultaneously accepted, but not that there exists a semantic model (a consistent assignment of truth-values to all propositions) that would make all the accepted propositions simultaneously true. On the given weak syntactic notion of consistency, the set $\{P, (P \rightarrow Q), \neg Q\}$, for instance, is consistent, since no proposition and its negation are both contained in it. On a stronger semantic notion of consistency, by contrast, the set is not consistent, since there exists no semantic model that would make all the propositions in this set simultaneously true. The conjunction of our weak syntactic notion of consistency and deductive closure entails the stronger semantic notion.

⁴ Systematicity is a undeniably a demanding condition. It requires that the collective judgment on a given proposition be dependent only on the individual assent/dissent pattern on that proposition, not on what the proposition is. One might weaken systematicity by allowing a dependency of the collective judgment not only on the individual assent/dissent pattern on the given proposition, but also on its syntactic structure. An interesting question for future research is the following. Define an equivalence relation \sim on the set X of all propositions -- for instance by stipulating $\phi \sim \psi$ if and only if ϕ and ψ have the same syntactic structure --, and ask whether the impossibility result of the

present paper persists, or fails to persist, if condition (S) is replaced with the following, weaker condition (S \sim).

SYSTEMATICITY WITH RESPECT TO \sim (S \sim). For any two propositions ϕ and ψ in X , where ϕ and ψ belong to the same equivalence class with respect to \sim , if every individual in N makes exactly the same judgment (assent/dissent) on ϕ as he or she makes on ψ , then the collective judgment (assent/dissent) on ϕ should also be the same as that on ψ , and the same pattern of dependence of collective judgments on individual ones should hold for all profiles in the domain of F .

Condition (S) is simply the special case of (S \sim) where all propositions in X belong to the same equivalence class with respect to \sim . We here restrict our consideration to condition (S) because it seems to us that the condition of systematicity is intuitively most compelling if it is interpreted as the requirement that *all* propositions be treated in an even-handed way. Propositionwise majority voting is a paradigm example of a judgment aggregation function satisfying systematicity in this strong form.

⁵ It is worth noting that, if even numbers of individuals are permitted and the domain of admissible profiles of personal sets of judgments is not restricted, propositionwise majoritary voting is already ruled out by the conditions of completeness and consistency of collective sets of judgments. For in cases of ties between P and $\neg P$ -- an equal number of individuals supporting P and $\neg P$ -- either both or neither will be supported collectively by propositionwise majority voting, depending on how ties are dealt with. At first sight, this might seem to suggest that our impossibility result is just an unsurprising consequence of the possibility of ties. However, the result is much more general and holds for any number of individuals, including an odd number, where there can be no ties if personal sets of judgments are complete and consistent.

⁶ This paper developed from joint discussion and exploration of themes raised in Pettit (2001b) but the formal proof of the impossibility result is the work of

Christian List. We wrote the paper when List was a Harsanyi Program Visitor at the Research School of Social Sciences, ANU, in March-Sept 2000. We wish to express our gratitude to Geoffrey Brennan, Campbell Brown, John Dryzek, Robert Goodin, Wlodek Rabinowicz and an anonymous reviewer for helpful comments and discussion. The paper was presented on a number of seminars, including a conference in honour of Isaac Levi, at Columbia University in October 2000. We were greatly helped by the comments of participants, in particular the follow-up comments given us by John Collins and Teddy Seidenfeld.

⁷ Note that the use of conjunction (\wedge) here is not essential, but that the use of other logical connectives would yield a similar result. Particularly, as the set of connectives $\{\neg, \wedge\}$ is expressively adequate, any logically possible proposition of the propositional calculus can be expressed as a proposition using \neg and \wedge as the only connectives.

⁸ We are indebted to Marek Kaminski for a suggestion on how to substantially simplify the proof.