

Groups with Minds of their Own

Philip Pettit

Introduction

There is a type of organisation found in certain collectivities that makes them into subjects in their own right, giving them a way of being minded that is starkly discontinuous with the mentality of their members. This claim in social ontology is strong enough to ground talk of such collectivities as entities that are psychologically autonomous and that constitute institutional persons. Yet, unlike some traditional doctrines (Runciman 1997), it does not spring from a rejection of common sense. This paper shows that the claim is supported by the implications of a distinctive social paradox — the discursive dilemma — and is consistent with a denial that our minds are subsumed in a higher form of Geist or in any variety of collective consciousness. Although the paper generates a rich, metaphysical brew, the ingredients it deploys all come from austere and sober analysis.

The paper is in six sections. In the first I introduce the doctrinal paradox, a predicament recently identified in jurisprudence, and in the second I explain how it generalises to constitute the discursive dilemma. In the third section I show that that dilemma is going to arise for any group or grouping — henceforth I shall just say, group — that espouses or avows purposes, and that such purposive collectivities are bound to resolve it by imposing the discipline of reason at the collective rather than the individual level. In the fourth and fifth sections I argue that groups of this kind — social integrates, as I call them — will constitute intentional and personal subjects. Then in the sixth and last section I look briefly at how we should think of the relationship between institutional persons of this kind and the natural persons who sustain them.

1. The doctrinal paradox

The discursive dilemma is a generalised version of the doctrinal paradox that has recently been identified in jurisprudence by Lewis Kornhauser and Lawrence Sager (Kornhauser and Sager 1986; Kornhauser 1992; Kornhauser 1992;

Kornhauser and Sager 1993. See too Chapman, 1998a, 1998b; Brennan 1999). This paradox arises when a multi-member court has to make a decision on the basis of received doctrine as to the considerations that ought to determine the resolution of a case: that is, on the basis of a conceptual sequencing of the matters to be decided (Chapman 1998a) . It consists in the fact that the standard practice whereby judges make their individual decisions on the case, and then aggregate their votes, can lead to a different result from that which would have ensued had they voted instead on whether the relevant considerations obtained, and let those votes dictate how the case should be resolved.

A good example of the doctrinal paradox is provided by this simple case where a three-judge court has to decide on a tort case. Under relevant legal doctrine let us suppose that the court has to judge the defendant liable if and only if it finds, first, that the defendant's negligence was causally responsible for the injury to the plaintiff and, second, that the defendant had a duty of care towards the plaintiff. Now imagine that the three judges, A, B and C vote as follows on those issues and on the doctrinally related matter of whether the defendant is indeed liable.

	Cause of harm?	Duty of care?	Liable?
A.	Yes	No	No
B.	No	Yes	No
C.	Yes	Yes	Yes

Matrix 1

There are two salient ways in which the court might in principle make its decision in a case like this. Let us suppose that each judge votes on each premise and on the conclusion, and does so in a perfectly rational manner. The judges might aggregate their votes in respect of the conclusion — the liability issue — and let the majority view on that issue determine their collective finding. Call this the conclusion-centred procedure. Under such a procedure, the defendant would go free, since there are two votes against liability. Or the judges might aggregate their votes on the individual premises — the causation and duty issues; let the majority view on each premise determine whether or not it is collectively endorsed; and let the conclusion be accepted — that the defendant is

liable — if and only if both premises are collectively endorsed. Call this the premise-centred procedure. Since each premise commands majority support, the defendant would be found liable under this procedure. The doctrinal paradox, as presented in the jurisprudential literature, consists in the fact that the two procedures described yield different outcomes.

Another simple example from the jurisprudential area is provided by a 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 the sort of contract that was allegedly in place. Now imagine that the three judges, A, B and C vote as follows on those issues and on the doctrinally related matter of whether the defendant is indeed liable.

	Contract?	Breach?	Liable?
A.	Yes	No	No
B.	No	Yes	No
C.	Yes	Yes	Yes

Matrix 2

In this case, as in the previous example, the judges might each conduct their own reasoning and then decide the case in a conclusion-centred way, by reference to the votes in the final column. Or they might decide the case in a premise-centred way by looking to the majority opinions in each of the first two columns and then letting those opinions decide the issue of liability. If they adopted the conclusion-centred approach, they would find for the defendant; if they took the premise-centred approach, then they would find against.

The paradox illustrated will arise wherever a majority in the group supports each of the premises, different majorities support different premises, and the intersection or overlap of those majorities is not itself a majority in the group. The fact that those in that overlap are not themselves a majority — in the cases considered there is only one judge, C, in the intersection — explains why there is only a minority in favour of the conclusion.¹

The doctrinal paradox is not confined to cases where a court has to make a decision by reference to a conjunction of premises. It can also arise in cases where the court has to make its decision by reference to a disjunction of considerations; that is, in cases where the support required for a positive conclusion is only that one or more of the premises be endorsed. This is unsurprising, of course, given that a disjunction of premises, p or q , is equivalent to the negation of a conjunction: $\text{not}(\text{not-}p \text{ and not-}q)$. Still, it may be worth illustrating the possibility.

Imagine that three judges have to make a decision on whether or not someone should be given a retrial; that a retrial is required either in the event of inadmissible evidence having been used previously or in the event of the appellants's having been forced to confess; and that the voting goes as follows among the judges (Kornhauser and Sager 1993, 40):

	Inadmissible evidence? Forced confession?		Retrial?
A.	Yes	No	Yes
B.	No	Yes	Yes
C.	No	No	No

Matrix 3

This case also illustrates a doctrinal paradox, since the conclusion-centred procedure will lead to giving the defendant a retrial and a premise-centred procedure will not: at least not, so long as majority voting is all that is required for the group to reject one of the premises (see Pettit 2001b).

2. The discursive dilemma

It should be clear that the doctrinal paradox will generalise in a number of dimensions, representing a possibility that may materialise with any number of decision-makers greater than two and with any number of premises greater than one, whether those premises be conjunctively or disjunctively organised. But there are other, perhaps less obvious ways in which it can be generalised also and I now look at three of these. These give us reason, as we shall see later, to speak of a discursive dilemma. I describe them respectively as the social generalisation, the diachronic generalisation, and the modus tollens generalisation.

The social generalisation

A paradox of the sort illustrated 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 will arise 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.

Consider an issue that might arise in a workplace, among the employees of a company: for simplicity, as we may assume, a company owned by the employees. The issue is whether to forego a pay-rise in order to spend the money thereby saved on introducing a set of workplace safety measures: say, measures to 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 letters A, B, and C represent the three employees and the ‘Yes’ or ‘No’ on any row represents the disposition of the relevant employee to admit or reject the corresponding premise or conclusion.

	Serious danger?	Effective measure?	Bearable loss?	Pay-sacrifice?
A.	Yes	No	Yes	No
B.	No	Yes	Yes	No
C.	Yes	Yes	No	No

Matrix 4

If this is the pattern in which the employees 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 are more 'Yes's' than 'No's' in each of the premise columns.

There are familiar practices of group deliberation and decision-making corresponding to the conclusion-centred and premise-centred options. Thus the group would go the conclusion-centred way if members entered into deliberation and dialogue and then each cast their personal vote on whether to endorse the pay-sacrifice or not; in that case the decision would be against the pay-sacrifice. The group would go the premise-centred way, on the other hand, if there was a chairperson who took a vote on each of the premises — say, a show of hands — and then let logic decide the outcome; in this case the decision would be in favour of the pay-sacrifice.

This example is stylised but should serve to indicate that the paradox is not confined to the domain in which legal doctrine dictates that certain judgments are to be made by reference to certain considerations. There are many social groups that have to make judgments on various issues and that routinely do so by reference to considerations that are privileged within the group.

One set of examples will be provided by the groups that are charged by an external authority with making certain decisions on the basis of designated considerations, and on that basis only. Instances of the category will be 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. With all such groups there is likely to be a problem as to whether the group should make its judgment on a certain issue in a premise-centred or conclusion-centred way; it will always be possible that those procedures will lead in different directions.

For a second set of examples consider those groups where it is a matter of internal aspiration that members find common grounds by which to justify whatever line they collectively take. Think of the political movement that has to work out a policy program; 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 have access to an antecedently agreed set of considerations on the basis of which to justify particular judgments. But their identification with one another will support a wish to reach agreement on such a set of reasons. To the extent that that wish gets to be satisfied, they will have to face the issue, sooner or later, as to whether they should make their decisions in a premise-centred or conclusion-centred way.

The diachronic generalisation

For all that has been said, however, the paradox may still seem unlikely to figure much in ordinary social life. The reason is that whereas the judges in a courtroom routinely have to make their judgments by reference to shared considerations, people in other social groups will often reach collective decisions on an incompletely theorised basis (Sunstein 1999). There will be a majority, perhaps even a consensus, in favour of a certain line on some issue but there will be no agreement among the parties to that majority or consensus on the reasons that support the line. The parties will each vote that line for reasons of their own — reasons related to their own interests or their own judgments of the common interest — and there will only be a partial overlap between the different considerations they each take into account. Thus there will be no possibility of their resorting to a premise-centred procedure, let alone any prospect of that procedure yielding a different result from the conclusion-centred alternative.

But sound as this consideration is, social groups will still have to deal routinely with the choice between these two procedures. In all of the examples so far considered, the premises and the conclusion are up for synchronic determination, whether at the individual or the collective level. Under the conclusion-centred procedure, each person has to make up their own mind on the reasons they are considering in premise position — assuming they do judge by reasons — and at the same time on the conclusion that those reasons support. Under the premise-centred procedure the group has to make up its mind on the reasons that are relevant by everyone's lights and at the same time on the conclusion that is to be derived from those premise-judgments. But the problem of choosing between such procedures may arise for a group in a diachronic as distinct from a synchronic way and is likely to arise much more generally on this basis.

Suppose that over a period of time a group makes a judgment on each of a set of issues, deciding them all by majority vote and perhaps deciding them on incompletely theorised grounds: different members of the group are moved by different considerations. Sooner or later such a group is bound to face an issue such that how it should judge on that issue is determined by the judgments it previously endorsed on other issues. And in such an event the group will face the old choice between adopting a conclusion-centred procedure and adopting a premise-centred one. The members may take a majority vote on the new issue facing them, running the risk of adopting a view that is inconsistent with the views that they previously espoused as a collectivity. Or they may allow the previously espoused views to dictate the view that they should take on this new issue.

The courts will often face diachronic examples of the problem illustrated as well as the synchronic examples that we considered; this will happen when previous judgments of the court dictate the judgment that it ought to make on an issue currently before it. But, more important for our purposes, even social groups that differ from the courts in routinely securing only incompletely theorised agreements will have to confront diachronic examples of the problem. They may escape the synchronic problem through not being capable of agreeing on common considerations by which different issues are to be judged. But that is no guarantee that they will be able to escape the problem as it arises in diachronic form.

The modus tollens generalisation

The third and last point to note in generalisation of the doctrinal paradox is that the options that we have been describing as the conclusion-centred procedure and the premise-centred procedure are not exhaustive of the alternatives available. The problem involved in the doctrinal paradox, even as it arises in legal and synchronic contexts, has a more general cast than the jurisprudential literature suggests.

The best way to see that the options are not exhaustive is to consider what a group may do if it finds that, relying on majority vote, it endorses each of a given set of premises while rejecting a conclusion that they support: say, deductively support. One grand option is for the collectivity to let the majority vote stand on each issue, thereby reflecting the views of its members on the different issues, while allowing the collective views to be inconsistent with one another. This approach, in effect, would vindicate the conclusion-centred procedure. But what now are the alternatives?

One possibility is for the group to ignore the majority vote on the conclusion, as in the premise-centred procedure, and to let the majority votes on the premises dictate the collective view on the conclusion. But another equally salient possibility, neglected as irrelevant in the legal context, is to ignore the majority vote on one of the premises, letting the majority votes on the other premises together with the majority vote on the conclusion dictate the collective view to be taken on that premise. The first possibility involves the collectivity practising modus ponens, the second has it practise modus tollens instead. These two options can be seen as different forms of a single grand option that stands exhaustively opposed to the first alternative described above. Where that alternative would have the collectivity reflect the individual views of its members on each issue, this second option would have the group ensure that the views collectively espoused across those issues are mutually consistent.

It should now be clear why I speak of a discursive dilemma rather than a doctrinal paradox. The problem arises because of the requirements of discourse as such, not just because of the demands of legal doctrine. And the problem represents a hard choice or dilemma, not anything that strictly deserves to be called a paradox.

The hard choice that a group in this dilemma faces is whether to let the views of the collectivity on any issue be fully responsive to the individual views of members, thereby running the risk of collective inconsistency; or whether to ensure that the views of the group are collectively rational, even where that means compromising responsiveness to the views of individual members on one or another issue. You can have individual responsiveness or collective rationality but you cannot have both — or at least you cannot have both for sure.

In arguing that the discursive dilemma presents groups with a hard choice, of course, I am assuming that they will not be happy to avoid that choice by insisting on voting by unanimity rather than majority, for example, since that would make them unable to come to agreement on many pressing questions. And equally I am assuming that collectivities will not simply refuse to draw out the implications of their views, avoiding inconsistency by avoiding deductive closure. But I say no more here on the general possibilities that arise in this area. Christian List and I have argued elsewhere for a relevant impossibility theorem (List and Pettit 2001).²

3. Resolving the dilemma by collectivising reason

Any groups that seek to make deliberative, reasoned judgments, then, face a dilemma. They may maximise responsiveness to individual views, running the risk of collectively endorsing inconsistent sets of propositions. Or they may impose the discipline of reason at the collective level, running the risk of collectively endorsing a conclusion that a majority of them — perhaps even all of them — individually reject. I show in this section that many groups respond to the dilemma by adopting the second alternative — by collectivising reason — and I go on to argue in the following two sections that groups which collectivise reason deserve ontological recognition as intentional and personal subjects.

Groups come in many different shapes and sizes (French 1984). Some are just unorganised collocations like the set of pedestrians on a given street, or the people who live in the same postal area. Some are sets related in other arbitrary ways, like those who have even telephone numbers or those who are first born to their mothers. And some are classes of people who share a common feature — say, accent or mannerism — that affects how others treat them but not necessarily how they behave

themselves. Yet other groups are united by a commonality, due to nature or artifice, that does affect how they behave themselves. It may affect how they behave towards one another, without leading them to do anything in common, as with linguistic groups, internet chat groups, and other enduring or episodic networks. Or it may also affect how they behave, as we say, to a shared purpose.

Purposive groups come themselves in a number of varieties (Stoljar 1973). They include organisations which have a specific function to discharge, such as museums, libraries, trusts, and states, as well as more episodic entities like the appointments committee or the jury or the commission of inquiry. And they also include groups which do not have any one specific function but which are associated with a characteristic goal, involving the outside world or the group's own members or perhaps a mix of both. Examples would include the political party, the trade union, and the business corporation, as well as the small group of colleagues involved in collaborative research and the set of friends arranging a joint holiday.

I argue in this section that purposive groups will almost inevitably confront examples of the discursive dilemma and that, short of resorting to deception, they will be under enormous pressure to collectivise reason: usually, though not inevitably, to collectivise reason by practising modus ponens — as in the premise-centred procedure — rather than modus tollens. In mounting this argument I shall speak as if every member of a purposive group participates equally with others in voting on what the group should do. I return to that assumption in the last section, where I try to show that the argument can survive variations in such detail.

My argument is in three parts. I argue, first, that a purposive collectivity will inevitably confront discursive dilemmas; second, that it will be under enormous pressure to collectivise reason in those dilemmas; and third, that in the general run of cases it will collectivise reason by following the premise-centred procedure.

The first part of the argument can be formulated in these steps.

1. Any collection of individuals who coordinate their actions around the pursuit of a common purpose — more on what this involves in the next section — will have to endorse judgments that dictate how they are to act; these will bear on the

opportunities available for action, the best available means of furthering their purpose, and so on.

2. The pursuit of such a common purpose will usually require explicit discussion and deliberation about the judgments the collectivity ought to endorse — it will not be like the activity of a tug-of-war team — so that over time the group will generate a history of judgments that it is on record as making.
3. Those past judgments will inevitably constrain the judgment that the group ought to make in various new cases; only one particular judgment in this or that case will be consistent — or coherent in some looser way — with the past judgments.
4. And so the group will find itself confronted with discursive dilemmas; it will be faced across time with sets of rationally connected issues such that it will have to choose between maximising responsiveness to the views of individual members and ensuring collective rationality.

This argument shows that discursive dilemmas of a diachronic sort are going to be more or less unavoidable for purposive groups but it is consistent, of course, with such groups also having to face synchronic dilemmas; I abstract from that possibility here. The second part of the argument goes on to show that any group of the kind envisaged will be pressured to impose the discipline of reason at the collective level. It involves a further three steps.

5. The group will not be an effective or credible promoter of its assumed purpose if it tolerates inconsistency or incoherence in its judgments across time; not all the actions shaped by those discordant judgments can advance, or be represented as advancing, one and the same purpose.
6. Every such group will need to be an effective promoter of its assumed purpose and will need to be able to present itself as an effective promoter of that purpose; it will lose any hold on members, or any respect among outsiders, if cannot do this.
7. And so every purposive group is bound to try to collectivise reason, achieving and acting on collective judgments that pass reason-related tests like consistency.

How will a purposive group be disposed to collectivise reason? We do not need to answer this question for purposes of the present argument. But it is worth noting that two plausible, further steps argue that such a group will generally,

though not of course inevitably, have to follow something like the premise-driven procedure illustrated in our earlier examples.

8. The group will be unable to present itself as an effective promoter of its purpose if it invariably seeks to establish consistency and coherence in the cases envisaged by renouncing one or other of its past commitments: if it never allows its present judgment to be dictated by past judgments; there will be no possibility of taking such a routinely inconstant entity seriously.
9. Thus, any such purposive collectivity must avoid automatic recourse to the revision of past commitments; it must show that those commitments are sufficiently robust for us to be able to expect that the group will frequently be guided by them in its future judgments.

The force of this three-part line of argument can be readily illustrated. Suppose that a political party announces in March, say on the basis of majority vote among its members, that it will not increase 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 increase government spending in other areas of policy or organisation. Should it allow a majority vote on that issue too?

If the party does allow a majority vote, then we know that even in the event of individual members being perfectly consistent across time, the vote may favour increasing government spending in other areas. Thus the party will face the hard choice between being responsive to the views of its individual members and ensuring the collective rationality of the views it endorses. The members may vote in the pattern of members A to C in the following matrix.

	Increase taxes?	Increase defence spending?	Increase other spending?
A.	No	Yes	No (reduce)
B.	No	No (reduce)	Yes
C.	Yes	Yes	Yes

Matrix 5

But the party cannot tolerate collective inconsistency, since that would make it a laughing-stock among its followers and in the electorate at large; it could no longer claim to be seriously committed to its alleged purpose. And so it must not allow its

judgments to be made in such a way that the discipline of reason is imposed only at the individual level; it has to ensure that that discipline is imposed at the collective level. In the ordinary run of things, the party will make its judgments after a premise-driven pattern, using a modus ponens pattern. It may occasionally revoke earlier judgments in order to be able in consistency to sustain a judgment that is supported by a majority. But it cannot make a general practice of this, on pain of again becoming a laughing-stock. It must frequently allow past judgments to serve as endorsed premises that dictate later commitments.

This argument with the political party is going to apply, quite obviously, to a large range of enduring and episodic collectivities. The argument does not rule out the possibility that those groups will occasionally adopt another course. They may choose to reject an earlier commitment in this or that case, for example, rather than revise their spontaneous judgment on the issue currently before them. Or they may even choose to live, overtly or covertly, with an inconsistency. But it is hard to see how they could generally fail in these regards and constitute effective or credible agents.

Instead of speaking of groups that collectivise reason in the manner of these collectivities I shall talk from now on of integrations of people, of integrated collectivities, and of social integrates. This way of speaking sounds a contrast with those groups that do not reason at all or that do not impose the discipline of reason at the collective level. These we naturally describe as aggregations of people, as aggregated collectivities or just as aggregates. I go on in the next two sections to argue that in an intuitive and important sense social integrates are going to be intentional and personal subjects. I continue to assume in this argument that members of social integrates all take an equal part in voting on what those collectivities should do; I come back to that assumption in the final section of the paper.

4. Social integrates are intentional subjects

Are integrations of people likely to constitute intentional subjects, displaying intentional states like beliefs and desires, judgments and intentions, and performing the actions that such states rationalise? In particular, are

integrations of people likely to constitute intentional subjects in their own right? Are we going to have to itemise them, side by side with their members — if you like, over and beyond their members — in any serious inventory of intentional subjects?

In a well-known discussion, Anthony Quinton (1975, 17) maintains not. He argues that to ascribe judgments, intentions and the like to social groups is just a way of ascribing them, in a summative way, to individuals in those groups.

We do, of course, speak freely of the mental properties and acts of a group in the way we do of individual people. Groups are said to have beliefs, emotions, and attitudes and to take decisions and make promises. But these ways of speaking are plainly metaphorical. To ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members. With such mental states as beliefs and attitudes, the ascriptions are of what I have called a summative kind. To say that the industrial working class is determined to resist anti-trade union laws is to say that all or most industrial workers are so minded.

The position adopted here by Quinton amounts to a straightforward eliminativism about collective intentional subjects. It suggests that only singular entities can constitute intentional subjects — for this reason it might also be called ‘singularism’ (Gilbert 1989, 12) — and that collectivities can be described as subjects ‘only by figment, and for the sake of brevity of discussion’ (Austin 1869, 364).

One reason why the position described amounts to eliminativism is this. If a collectivity can be said to form a certain belief or desire, a certain judgment or intention, so far as all or most of its members do, then it would be misleading to say that it constituted an intentional subject over and beyond its members. Asked to say how many such subjects were present in a certain domain it would be quite arbitrary to count the individuals there, and then to count the collectivity also. We might as well count as subjects, not just the total set of people there, but also every subset in which majority or unanimous attitudes give us a basis on which to ascribe corresponding attitudes to that collection of people.

This criticism suggests that Quinton tells too simple a story about the attitudes that we expect to find on the part of individuals of whom we say that they collectively judge or intend something. More recent work on the conditions that might lead us to ascribe such joint attitudes, and to posit collective subjects, has stressed the fact that we usually expect a complex web of mutual awareness on the part of individuals involved (Gilbert 1989; Searle 1995; Tuomela 1995; Bratman 1999). Thus Michael Bratman (1999) argues that you and I will have a shared intention to do something just in case a) you intend that we do it and I intend that we do it; b) we each intend that we do it because a) holds; and c) those clauses are matters of which we are each aware, each aware that we are each aware, and so on in the usual hierarchy of mutual knowledge.

Suppose we complicate the Quinton story in some such pattern, adopting one of these mutual-awareness analyses. Will that undercut his eliminativism, giving us reason to think that apart from singular subjects there are also collective ones? It will certainly evade the criticism just made, for it will make it much harder than Quinton does for a collection of individuals to deserve to be described as having certain mental properties. But it will not avoid another problem. It will not ensure that a collectivity displays the sort of rationality that we expect in the performance of any system we would describe as an intentional subject. So at any rate I shall argue.

What sort of rationality do we expect in an intentional subject? By a line of argument that has been widely endorsed in recent philosophical thought, a system will count as an intentional subject only if it preserves intentional attitudes over time and forms, unforms and acts on those attitudes — at least within intuitively feasible limits and under intuitively favourable conditions — in a rationally permissible manner: in a phrase, only if it displays a certain rational unity (Pettit 1993, Ch.1). If the system believes that *p* and comes across evidence that not *p*, it must tend to unform that belief. If the system believes that *p* and learns that if *p* then *q*, it must come to form the belief that *q* or to unform one of the other beliefs. If the system desires that *p*, believes that by *X*-ing it can bring it about that *p*, and believes that other things are equal, then it must tend to *X*. And so on.

Even if we introduce the sort of complexity postulated in mutual-awareness stories about collective subjects, that will not guarantee that those subjects have the rational unity associated with intentionality. Those stories are all consistent with the collectivity's acting by conventions that allow rational disunity. The convention established in the mutual awareness of members may ordain, for example, that the collectivity shall be deemed to judge or intend whatever a majority of members vote for its judging or intending at that time. And we know from discussion of the discursive dilemma that if such a convention obtains — if the attitudes of the collectivity are required to be continuous in that majoritarian way with the current votes of members — then the collectivity may be guilty of grievous irrationality over time. It may be as wayward in the postures it assumes as the most casual aggregate of individuals; it may fail to materialise as anything that deserves to be taken as an intentional subject in its own right.

In order for a collectivity to count as an intentional subject, not only must there be a basis in the behaviour of participating members for ascribing judgments and intentions and such attitudes to the collective; that is the point on which the mutual-awareness literature rightly insists. There must also be a basis for thinking of the collectivity as a subject that is rationally unified in such a way that, within feasible limits and under favourable conditions, we can expect it to live up to the constraints of rationality; we can expect it to enter and exit states of belief and desire, judgment and intention, in a way that makes rational sense and we can expect it to perform in action as those states require. Indeed, were there a basis for ascribing such states to a collectivity, and a basis for expecting this sort of rational unity, then it is hard to see any reason why we should deny that the collectivity was an intentional subject in its own right.

How to secure the dual basis that is necessary for a collectivity to be an intentional subject? The argument of the last section suggests a salient recipe. By ensuring that the collectivity represents an integration of individuals, not just a casual aggregate. Specifically, by ensuring, first, that the collectivity has a shared purpose and forms the judgments and intentions associated with pursuit of that

purpose; and second, that it collectivises reason in forming those judgments and intentions.

I said and say nothing on what it is for a collectivity to form and have a shared purpose, or to form and have certain judgments and intentions. Presumably that can be analysed on something like the lines explored in the mutual-awareness approach; it has to do, plausibly, with the conventions, and the associated structures of common knowledge, that prevail in the collectivity. Assuming that there is an established, conventional sense in which a collectivity has a shared purpose, and forms associated judgments and intentions, the fact that it collectivises reason in the course of that enterprise — the fact that it is a social integrate — means that it will display precisely the sort of rational unity required of an intentional subject. Let the collectivity have made certain judgments and formed certain intentions in the past. And now imagine that it faces a theoretical or practical issue where those judgments and intentions rationally require a particular response. We can rely on the integrated collectivity to respond as those intentional states rationally require, or to make rationally permissible adjustments that undercut the requirements. Or at least we can rely on it to do this under intuitively favourable conditions, and within intuitively feasible limits.

The integrated collectivity has common purposes, and forms associated judgments and intentions, unlike the collections envisaged in Quinton's account. And the integrated collectivity can be relied upon to achieve a rational unity in the judgments and intentions endorsed, unlike the group that meets only the mutual-awareness conditions for forming collective attitudes. It satisfies the dual basis that is necessary for a collectivity to count as an intentional subject. But is the satisfaction of these two conditions sufficient as well as necessary for the integrated collectivity to count as an intentional subject, in particular an intentional subject that is distinct from the individual subjects who make it up?

If we are to recognise the integrated collectivity as an intentional subject, then we must admit of course that it is a subject of an unusual kind. It does not have its own faculties of perception or memory, for example, though it may be able to register and endorse facts perceived or remembered by others: in

particular, by its own members. Under our characterisation it is incapable of forming degrees of belief and desire in the ordinary fashion of animal subjects; its beliefs are recorded as on-off judgments, its desires as on-off intentions. And the judgments and intentions that it forms are typically restricted to the narrow domain engaged by the particular purposes that its members share. Notwithstanding these features, however, I think that it is reasonable, even compulsory, to think of the integrated collectivity as an intentional subject.

The basis for this claim is that the integrated collectivity, as characterised, is going to display all the functional marks of an intentional subject and that there is no reason to discount those marks as mere appearances. Within relevant domains it will generally act in a manner that is rationalised by independently discernible representations and goals; and within relevant domains it will generally form and uniform those representations in a manner that is rationalised by the evidence that we take to be at its disposal. In particular, it will manifest this sort of functional organisation, not just at a time, but over time; it will display the degree of constancy as well as the degree of coherence that we expect in any intentional subject. Why would anyone deny that an entity that displays the functional marks of an intentional subject in this manner is not really an subject of that kind?

One ground might be that intentionality requires, not just a certain form of organisation, but also the realisation of that form in inherently mental material, whatever that is thought to be. Few would endorse this consideration among contemporary thinkers, however, since there appears to be nothing inherently mental about the biological material out of which our individual minds are fashioned (but see Searle 1983). Another ground for the denial might be that the functional marks of intentional subjectivity have to come about as a result of the subject's internal organisation, and not in virtue of some form of remote control or advance rigging (Jackson 1992). But this is hardly relevant to the integrated collectivity, since its judgments and intentions are clearly formed in the required, internal fashion. Still another sort of ground for denying that functional organisation is sufficient for being an intentional subject is that something more is required — say, natural selection or individual training (Millikan 1984;

Papineau 1987; Dretske 1988) — for the attitudes of the subject to have determinate contents. This is not relevant in the case of the integrated collectivity, however, since the contents of its judgments and intentions will inherit determinacy from the presumptively determinate words that are used by its members to express those contents.

The usual grounds for driving a wedge between functionally behaving like an intentional subject and actually being an intentional subject are unlikely, as this quick survey shows, to cause a problem with the integrated collectivity. If further grounds for making such a separation between appearance and reality are lacking, then we have every reason to treat the integrated collectivity as an intentional subject. And such grounds, so far as I can see, are indeed lacking. I can think of only one other consideration that might be invoked against counting integrated collectivities as intentional subjects and it does not raise a serious problem.

The consideration is that if we treat integrated collectivities as intentional subjects, then we may be involved in a sort of double-counting. We will be counting the individual members of the collectivity as intentional subjects. And then we will be going on to say that apart from those members, there is a further subject present too: the collectivity that they compose. But I do not think that this makes for an objection. The integrated collectivity will not be distinct from its individual members, in the sense that it will not be capable of existing in the absence of such members. But it will be distinct in the sense of being a centre for the formation of attitudes that are capable of being quite discontinuous from the attitudes of the members. This is one of the lessons of the discursive dilemma.

Consider the case of the worker-owners who have to decide on whether to forego a pay-rise in order to purchase a device for guarding against the danger of electrocution. Imagine that they cast their votes after the pattern illustrated in Matrix 4 and that they follow the premise-centred procedure in determining what to think about the issue. In such a case the group will form a judgment on the question of the pay-sacrifice which is directly in conflict with the unanimous vote of its members. It will form a judgment that is in the starkest possible discontinuity with the corresponding judgments of its members.

As the point applies to judgment, so it naturally extends to intention. The collectivity of workers that makes a judgment in favour of the pay-sacrifice will be firmly disposed to act accordingly, under the procedure it adopts, and in that sense it will form a corresponding intention. Thus the chairperson will be entitled by the premise-driven procedure to announce on the basis of the premise-votes: 'Colleagues, our intention is fixed: we will forego the pay-rise'. But at the moment where the intention of the integrated group is thereby fixed, no one member will intend that the group act in that way, or that he or she play their part in the group's acting in that way. Such individual intentions will follow on the formation of the group intention, of course, since the group can only act through the actions of its members. But they are not the stuff out of which the group intention is constructed; on the contrary, they are effects that the formation of the group intention plays a role in bringing about.

These discontinuities between collective judgments and intentions, on the one hand, and the judgments and intentions of members, on the other, make vivid the sense in which a social integrate is an intentional subject that is distinct from its members. They represent the cost that must be paid if a collectivity is to achieve the rational unity that we expect in any intentional subject. Rational unity is a constraint that binds the attitudes of the collectivity at any time and across different times, and the satisfaction of that constraint means that those attitudes cannot be smoothly continuous with the corresponding attitudes of members.

In arguing that a social integrate is an intentional subject that is distinct from its members — that exists over and beyond its members — I hasten to add that I am not postulating any ontological mystery. The argument is consistent with the supervenience claim that if we replicate how things are with and between individuals in a collectivity — in particular, replicate their individual judgments and their individual dispositions to accept a certain procedure — then we will replicate all the collective judgments and intentions that the group makes. Collective judgments and intentions may be discontinuous with what happens at the individual level but they cannot vary independently of what happens there; they do not constitute an ontologically emergent realm.³

5. Social integrates are institutional persons

This discontinuity between an integrated collectivity and its members, and the fact that such a collectivity can constitute a distinct intentional subject, is quite surprising. But there is more to come. For it turns out that the way in which the judgments and intentions of social integrates are formed and policed forces us to think of those collectivities as institutional persons. It leads us to see that like individual human beings, and unlike non-human animals, they display everything that is strictly necessary in personal as distinct from just intentional subjects.

What distinguishes personal from merely intentional subjects? As I assumed in the previous discussion that intentional subjects have to display a certain rational unity, so I make a parallel assumption in discussing this question. I assume that whereas intentional subjects must have intentional states and perform associated actions in a way that satisfies rational unity — whether or not they are aware of doing so — persons must be capable of being held to that ideal; they must be such that they can be held responsible for failures to unify their intentional states and actions in a rational way (Rovane 1987; Pettit 2001a, Ch. 4). Rational unity is a constraint that intentional systems must be designed to fulfill, if only at subpersonal, unconscious levels. Rational unification is a project for which persons must be taken to assume responsibility, at least on a case by case basis.

The commitment that persons make to rational unification, according to this account, means that persons don't just possess intentional states, and perform corresponding actions. They also avow those states and actions, acknowledging them as their own. And, avowing them, they hold themselves open to criticism in the event of not proving to live up to them: not proving to satisfy rational unity in their regard. Let a person avow a belief that *p* and a belief that if *p* then *q*, for example, and we can expect them to form and avow the belief that *q*. Or if they fail to do so, then we can expect them to have a justification or an excuse to offer. The justification may be that they had a change of mind in respect of '*p*' or 'if *p* then *q*', the excuse that the conditions under which they were operating made it difficult to think straight.

The assumption that persons are marked off from ordinary intentional subjects — say, non-human animals — by the commitment to rational unification makes for a rich conception of personhood. But for that very reason it will hardly be contested in the present context, for the richness of the account should make it harder rather than easier to argue that integrations of people count as persons. In any case I will say nothing more in its defence here. I shall take as persons those intentional agents who can avow their intentional states and the actions they perform in words — or in signs of some other sort — and who can then be held to the associated expectations. We may describe as persons those human beings who do not yet have this capacity, who no longer have it, or who do not have it at all. But that usage is readily seen as an extension based on the fact that they are of a kind — that is, of a species — with creatures who are persons in that strict sense.

Assuming that persons are intentional agents who make and can be held to avowals, what are we to say of integrated groups? I have no hesitation in arguing that this means that they are institutional persons, not just institutional subjects or agents (Rovane 1997 argues a similar line). Integrated collectivities bind themselves to the discipline of reason at the collective level, and that means that they are open to criticism in the event of not achieving rational unity in relevant regards. They avow judgments, intentions and actions and prove able to be held responsible for failures to achieve consistency and other such ideals in their associated performance. They are subjects which can be treated as properly conversable interlocutors (Pettit and Smith 1996).

Social integrates contrast in this respect with any groups that do not impose the discipline of reason at the collective level. Collectivities of this aggregate kind will not be answerable in the same way to words previously authorised or deeds previously performed. And that will be so, no matter how tight we make the mutual-awareness constraints on when they can be said to authorise words or perform deeds. It will always be possible for such an aggregate to vote in favour of a judgment or an intention or an action that is out of kilter with earlier commitments, and to do so without being exposed to legitimate criticism. Opinion poll research may tell us that the populace as a

whole supports cutting taxes, increasing defence expenditure and also increasing other expenditure. Since all the individuals involved may hold consistent views, that finding that will not give us reason to criticise the populace for holding such opinions, as we might criticise a political party for doing so. For even if it is taken as an intentional subject, the populace cannot be expected to police itself for the rational unity of the things it believes and then be held to that expectation. It does not constitute a person in the relevant sense and it contrasts in that regard with the political party.

Whenever we speak of persons we think it is appropriate to speak of selves. We expect that persons will think of themselves in the first person and be able to self-ascribe beliefs and desires and actions by the use of an indexical expression like 'I' or 'my', 'me' or 'mine'. This association between being a person and thinking in self-ascriptive terms is borne out under the characterisation of persons adopted here. If a person is to avow certain states and actions, and assume responsibility for achieving rational unity in their regard, then those states and actions are bound to have a distinctive salience in their experience. Individual subjects are bound to see them — by contrast with the states and actions of others — as matters of what I believe, what I desire, what I do, and so on (Pettit 2001a, Ch.4).

Why must the personal point of view have this indexical, first-personal character? Why must I as a conversable subject be aware of myself in this indexical way, rather than just under a name, say as PP? A well-known line of argument provides the answer (Perry 1979; Burge 1998). Were I to conceive of myself under a name, as PP, then there would always be a deliberative gap between my thinking that PP believes both that *p* and that '*p*' entails '*q*' and my actually adjusting beliefs — say, in response to conversational challenge — by coming to believe that *q* or by giving up one of the other beliefs. For why should my beliefs about PP's beliefs have any reason-mediated effect on what I believe and assert, short of my believing that I am PP? And if I can think that I am PP, of course, then I do think of myself in the first person, not just under a name.

So far as integrated collectivities operate on the same lines as individual persons, they will also have this capacity to think in first person terms. From the

standpoint of those in an integrated collectivity the words defended in the past, for example, will stand out from any words emanating from elsewhere as words that bind and commit them. Specifically, they will stand out for those of us in the collectivity as words that ‘we’ as a plural subject maintain. The argument in the singular case for why I as a person must conceive of my attitudes as matters of what I think applies in the plural case too, showing that we, the members of an integrated collectivity, must think of the group’s attitudes as matters of what we think.

The members of a social integrate, S, will face the same deliberative gap as that which appeared in the singular case, if they conceive of the existing commitments of the group just as those that holds. Suppose that we in that group recognise both that p and that the truth of ‘p’ entails the truth of ‘q’. That will not lead us as a group to judge that q, unless we make the extra judgment that we are S. And if we do make that judgment then of course we do think of ourselves in the first person plural. As members of the integrated group, we are possessed of a personal point of view and it is marked out by this indexical usage.

The emphasis on the importance of ‘we’ connects with the insistence by writers like Margaret Gilbert (1989), John Searle (1995) and Annette Baier (1997) that there is no possibility of analysing we-talk in I-talk, or indeed in impersonal talk of what named individuals do (see too Tuomela 1995, 183). The obstacle to reducing talk of ‘we’ to talk of ‘I’ will be just the obstacle that stands in the way of reducing indexical talk of what I think and do to non-indexical talk of what PP thinks and does. As there is a personal perspective that is available only with talk of ‘I’, so there is a personal perspective that becomes available only with talk of ‘we’.

The autonomy of ‘we’ talk that has to obtain under our account of what it is for a collectivity to be integrated nicely emphasises the significance of the claim that such collectivities are personal as well as intentional agents. Not only do social integrates have a rational unity that constrains their performance over time and that makes them distinct from their own members. The rational unity they display is one that they themselves police and implement in the fashion of creatures whom we can hold responsible: creatures who count as persons

(McGeer and Pettit 2001). They are rationally unifying as well as rationally unified subjects and the enterprise of unification in which they are involved forces them to think in the manner of a self. It makes it natural and indispensable for members to resort to a distinctively proprietary use of ‘we’, ‘us’ and ‘ours’.

Once again, I should say, there is no ontological mystery in any of this: no suggestion of reality sundering into distinct realms, for example, indexical and non-indexical. If we fix the way the world is in impersonal, non-indexical terms, then we will have fixed all the indexical truths — in particular all the I-truths and all the we-truths — as well. Indexical truths supervene on non-indexical, since the same indexical sentences will be true at the same locations of utterance in impersonally indiscernible worlds (Jackson 1998; Pettit 2000). But this sort of fixing — this ontological reducibility — is quite consistent with the perspective of non-indexical talk failing to register things in the way required for singular conversability, and indeed with the perspective of I-talk failing to register things in the way required for plural conversability. Such idioms may fail to be intertranslatable, and yet not direct us to independent realms of reality.

6. Natural and institutional persons

The claim just defended is that social integrates have to be regarded as persons, on a par with individual human beings. But it is consistent, of course, with acknowledging that such institutional persons differ from natural persons in as many ways as they resemble them. As we saw earlier, institutional persons are not centres of perception or memory or sentience, or even of degrees of belief and desire. Institutional persons form their collective minds only on a restricted range of matters, to do with whatever purpose they are organised to advance. And institutional persons are artificial creatures whose responses may be governed by reason, not in the spontaneous manner that is characteristic of individual human beings, but only in a painstaking fashion. Their reasoning may be as tortuous as that of the impaired human being who has to work out reflectively, case by case, that in virtue of believing that *p* and that if *p* then *q*, he or she ought also to believe that *q*. Integrated collectivities are persons in virtue of being conversable and responsible centres of judgment, intention and action. But they are persons of a bloodless, bounded and crudely robotic variety.

Even granted this, however, there are still important questions as to how institutional and natural persons relate to one another; in particular, how institutional persons and the members who constitute them relate to one another. I address two such questions in this brief conclusion.

The first is a more or less straightforward question as to the institutional profiles that members have to assume so far as they constitute a single collective person. Throughout this paper I have been assuming that even if just one individual has to act on behalf of a collective, the members are all equal participants in the formation of the collective's judgments and intentions, having equal voting power with others. But this is unrealistic as an assumption about most real-world collectives, and the first question is how far membership is consistent with the absence of such voting power.

There are two ways in which individuals may be said to endorse a collective procedure or outcome. First, by actively voting in its favour; and second, by having a capacity for exit or contestation or something of the kind — this, as a matter of common awareness — but not exercising that power. Although active voting is the most obvious mode of endorsement, it must also be possible for people to endorse a collective pattern in the second, virtual mode. The members of a collectivity cannot vote on every procedure that they are to follow, on pain of infinite regress; the problem is akin to that which would be involved in trying to endorse as an explicit premise every principle of inference deployed in an argument (Carroll 1895). If regress is to be avoided, therefore, then some of the procedures followed by the members of a collectivity must be followed without the endorsement of an explicit vote and just on the basis that that is how things are done among members, and done without contestation.

But if all the members of a group must endorse some procedures in a virtual way — that is, by not exercising a power of exit or contestation or whatever — then it is clearly possible that on many matters of procedure, and on many outcomes, some members will play an active voting part, while others are involved only in that virtual manner. And this is how it is, obviously, with most integrated collectivities. Such collectivities sometimes involve all of their members in deliberation on every decision. But more often they stage their

decisions so that the full assembly only votes on general matters, delegating others to smaller bodies and to officers of the group. Or they may involve a membership that is largely passive, with most being involved in official decisions only to the extent of needing to be pacified. Or they may be articulated into subunits that are each passive in relation to one another. And so on: the details need not concern us.

The second question raised by our discussion bears on how natural and institutional persons relate to one another within the psychology of a given member. Suppose that someone is faced with a decision on which they as a natural person tend to go one way, while an institutional person of which they are a member — perhaps the relevant, executive member — would tend to go another. What is to happen in such a case? For all that we have said, it might be that the psychology of the individual is taken over, willy nilly, either by the natural person associated with it or by the institutional person. Or it might be that which person is to be present in that psychology is determined, at least ideally, by considerations that the two persons can debate — debate within the same head, as it were — and reach agreement on. Or it might be that the natural person is always primary and has the task of deciding whether to act in their own name — in their own interests, perhaps, or according to their own values — or in the name of the collective. The first model is clearly crazy, suggesting that persons take over psychologies in the way demons are said to assume possession of souls. But which of the other two models is the more plausible?

My own inclination is to go for the last alternative, giving priority to natural persons. I reject the picture according to which persons, natural and institutional, are of more or less the same standing and have equal presumptive claims in the sort of case envisaged on the resources of the member's psychology (Rovane 1997 supports this image). I hold that natural persons have an inescapable priority and that in this kind of case it will be up to the natural person to decide whether or not to cede place to the institutional, acting in furtherance of the collective goal and in neglect of his or her own priorities.

There are a couple of reasons why I hold by this image rather than the other. One is that it fits well with the intentional manner in which, as it seems,

natural persons go about constituting and enacting institutional agents. Natural persons are in intentional control of whether they enter or exit most of the collectives to which they belong. And when they act on behalf of a collective they are reinforced in their identity as natural persons, and the intentional control they have as natural persons, by the way others relate to them; others call on Jones to do what the collective requires of them, others congratulate Jones for doing his or her bit, and so on.

Another reason for preferring my model is that there are cases where it is going to be quite misleading to think of two persons, one natural and the other institutional, debating within a single head as to who should be the one to prevail. That model may apply when the reasons that they take into account are agent-neutral considerations to do with what is for the best overall but it will be unrealistic where each person has an agent-relative reason — say, one to do with personal prospects or commitments or allegiances — for wanting to go their preferred way. When ordinary people diverge in that way, then reason runs out and they may have to compete in some non-deliberative manner — or toss a coin — to determine who wins. We cannot envisage a natural and an institutional person competing in that way within the same head.⁴

It is sometimes said that before we know what it is rational for a human being to do, we need to be told which identity that agent is enacting; in particular we need to be told whether they are acting in their own name or the name of a collectivity (Hurley 1989). Thus Elizabeth Anderson (2000) defends ‘The Priority of Identity to Rational Principle: what principle of choice it is rational to act on depends on a prior determination of personal identity, of who one is’. The line just taken suggests that this is not so. The natural person is the ultimate centre of action and if it is rational for a human being to act in the name of a collectivity — that is, rational in the sense of maximising relevant preferences — then it is rational in terms of the natural person’s preferences.

Conclusion

In maintaining points of the kind defended in this paper, we make contact with the tradition that the nineteenth century German historian, Otto von Gierke,

sought to track and to revitalise: the tradition of emphasising the institutional personality of many groups and the significance of such personality for legal, political and social theory (Hager 1989; Runciman 1997; McLean 1999). This tradition is deeply organicist in its imagery, and led adherents to speak for example of ‘the pulsation of a common purpose which surges, as it were, from above, into the mind and behaviour of members of any true group’ (Barker 1950, 61). But the organic, often overblown metaphors should not be allowed to discredit the tradition. The points they were designed to emphasise are perfectly sensible observations of the kind that our analysis of integrated groups supports.

I have argued elsewhere that consistently with being individualistic about the relation between human beings and the social regularities under which they operate — consistently with thinking that social regularities do not compromise individual agency — we may oppose the atomism that insists on the coherence of the solitary thinker; we may argue that individuals depend non-causally on one another for having the capacity to think (Pettit 1993). What we have seen in this paper is that consistently with being individualistic we may also oppose the singularism that insists on the primacy of the isolated agent and claims that we can describe collectivities as persons only in a secondary sense.

Individualism insists on the supervenience claim that if we replicate how things are with and between individuals, then we will replicate all the social realities that obtain in their midst: there are no social properties or powers that will be left out (Macdonald and Pettit 1981; Currie 1984; Pettit 1993). But this insistence on the supervenience of the social in relation to the individual is quite consistent with emphasising that the entities which individuals compose can assume a life of their own, deserving the attribution of discontinuous judgments and intentions and displaying all the qualities expected in personal agents.

The world of living organisms did not cease to be interesting when scientists dismissed the conceit of a vis vitalis. And neither should the world of social groups cease to be interesting, just because we choose to exorcise the spectre of a vis socialis. On the contrary, the recognition that the realm of collectivities is an artefact of human hands should excite the sociological, the political and the historical imagination. The sociological, because we badly need

general models of how collectivities can be created and sustained (Coleman 1974). The political, because we need to develop criteria for assessing the performance of collectivities and proposals for containing their power (Hager 1989). And the historical, because we have only the sketchiest understanding of how the most important collectivities in our lives emerged and stabilised as integrated agents (Skinner 1989).⁵

References

- Anderson, E. (April 2000). Unstrapping the Straitjacket of 'Preference': A Comment on Amartya Sen's Contributions to Philosophy and Economics. American Philosophical Association Meeting, Central Division, Chicago.
- Austin, J. (1869). Lectures on jurisprudence, or the philosophy of positive law. London.
- Baier, A. C. (1997). The Commons of the Mind. Chicago, Open Court.
- Barker, E. (1950). Introduction. Gierke, Natural Law and the Theory of Society. E. Barker. Cambridge, Cambridge University Press.
- Bratman, M. (1999). Faces of Intention: Selected Essays on Intention and Agency. Cambridge, Cambridge University Press.
- Brennan, G. (1999). Collective Irrationality and Belief. Canberra, Research School of Social Sciences, Australian National University.
- Burge, T. (1998). Reason and the First Person. Knowing Our Own Minds. B. S. C.Wright, and C.Macdonald. Oxford, Oxford University Press.
- Carroll, L. (1895). "What the Tortoise said to Achilles." Mind 4: 278-80.
- 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.
- Coleman, J. (1974). Power and the Structure of Society. New York, Norton.
- Currie, G. (1984). "Individualism and Global Supervenience." British Journal for the Philosophy of Science 35: 345-358.
- Dretske, F. (1988). Explaining Behavior. Cambridge, Mass., MIT Press.
- French, P. A. (1984). Collective and Corporate Responsibility. New York, Columbia University Press.
- Gilbert, M. (1989). On Social Facts. Princeton, N.J., Princeton University Press.

- Hager, M. M. (1989). "Bodies Politic: The Progressive History of Organizational "Real Entity" Theory." University of Pittsburgh Law Review 50: 575-654.
- Hurley, S. (1989). Natural Reasons. New York, Oxford University Press.
- Jackson, F. (1992). Block's Challenge. Ontology, Causality, and Mind: Essays on the Philosophy of David Armstrong. K. Campbell, J. Bacon and L. Rhinehart. Cambridge, Cambridge University Press.
- Jackson, F. (1998). From Metaphysics to Ethics: A Defence of Conceptual Analysis. Oxford, Oxford 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.
- List, C. and P. Pettit (2001). The Aggregation of Sets of Judgments: An Impossibility Result. Canberra, Australian National University.
- Macdonald, G. and P. Pettit (1981). Semantics and Social Science. London, Routledge & Kegan Paul.
- McGeer, V. and P. Pettit (2001). "The Self-regulating Mind." Language and Communication.
- McLean, J. (1999). "Personality and Public Law Doctrine." 49: 123-49.
- Meijers, A. (1994). Speech Acts, Communication and Collective Intentionality: Beyond Searle's Individualism. Utrecht, de Jonge.
- Millikan, R. (1984). Language, Thought and Other Biological Categories. Cambridge, Mass., MIT Press.
- Papineau, D. (1987). Reality and Representation. Oxford, Blackwell.

- Perry, J. (1979). "The Essential Indexical." Nous 13: 3-21.
- Pettit, P. (1993). The Common Mind: An Essay on Psychology, Society and Politics, paperback edition 1996. New York, Oxford University Press.
- Pettit, P. (2000). A Sensible Perspectivism. Dealing with Diversity. M.Baghramian and A.Dunlop. London, Routledge.
- Pettit, P. (2001). A Theory of Freedom: From the Psychology to the Politics of Agency. Cambridge, Polity.
- Pettit, P. and M. Smith (1996). "Freedom in Belief and Desire." Journal of Philosophy 93: 429-49.
- Quinton, A. (1975). "Social Objects." Proceedings of the Aristotelian Society 75.
- Rovane, C. (1997). The Bounds of Agency: An Essay in Revisionary Metaphysics. Princeton, NJ, Princeton University Press.
- Runciman, D. (1997). Pluralism and the Personality of the State. Cambridge, Cambridge University Press.
- Searle, J. (1995). The Construction of Social Reality. New York, Free Press.
- Searle, J. R. (1983). Intentionality. Cambridge, Cambridge University Press.
- Skinner, Q. (1989). The State. Political Innovation and Conceptual Change. T. Ball, J. Farr and R. Hansen. Cambridge, Cambridge University Press.
- Stoljar, S. J. (1973). Groups and Entities: An Inquiry into Corporate Theory. Canberra, Australian National University Press.
- Sunstein, C. (1999). One Case At A Time. Cambridge, Mass, Harvard University Press.
- Tuomela, R. (1995). The Importance of Us. Stanford, CA, Stanford University Press.

Philip.Pettit@anu.edu.au,
Research School of Social Sciences,
Australian National University,
Canberra, ACT 0200.

¹ The structure involved is this:

-
- a. there is a conclusion to be decided among the judges by reference to a conjunction of independent or separable premises — the conclusion will be endorsed if relevant premises are endorsed, and otherwise it will be rejected;
 - b. each judge forms a judgment on each of the premises and a corresponding judgment on the conclusion;
 - c. each of the premises is supported by a majority of judges but those majorities do not coincide with one another;
 - d. the intersection of those majorities will support the conclusion, and the others reject it, in view of a; and
 - e. the intersection of the majorities is not itself a majority; in our examples only one judge out of the three is in that intersection.

² Let the views of certain individuals on a rationally connected set of issues be rationally satisfactory in the sense of being consistent, complete and deductively closed. The impossibility theorem shows that any procedure whereby an equally satisfactory set of views may be derived from the individual views must fail in one of the following regards. It must be incapable of working with some profiles of individual view. Or it must fail to treat some individual or some issue even-handedly: roughly, it must let some individual or individuals be treated as less important than another — at the limit, the other may be given the status of a dictator — or it must downgrade some issue in the sense of letting the collective view on that issue be determined, not by majority vote, by the collective views on other issues.

³ There are other ontological questions that I do not address here. One is the issue of whether a group at any time is constituted in some sense by the individuals involved or is identical with the fusion of those individuals. This parallels the familiar sort of question raised about whether a statue is constituted by the body of clay used in its manufacture or whether it is identical with that body of clay. Different positions may be taken on this question, consistently with the claims made in the text.

⁴ This line of thought might be blocked by a consequentialist argument to the effect that all such divergences have to be judged ultimately by reference to

agent-neutral considerations. But it would be strange to tie one's view of the relationship between natural and institutional persons to a consequentialist commitment. And in any case it is possible for consequentialists to argue that it is often best for people — best in agent-neutral terms — to think and even compete in agent-relative ways.

⁵ I was greatly helped in developing this paper by conversations with Chandran Kukathas, Christian List and Victoria McGeer, as well as by comments from two anonymous referees. I am indebted to the discussion it received at a number of venues: the Summer Institute on 'Social Ontology after The Common Mind', held in Erasmus University, Rotterdam, July 2000, where my commentators were Ron Mallon and Rafal Wierzschoslawski; and at seminars in Columbia and Yale. I am also grateful for a useful set of comments from Elizabeth Anderson and for discussions of the priority of identity principle with Akeel Bilgrami and Carol Rovane.