

Group Agency and Supervenience

Christian List and Philip Pettit¹

Abstract. Can groups be rational agents over and above their individual members? We argue that group agents are distinguished by their capacity to mimic the way in which individual agents act and that this capacity must ‘supervene’ on the group members’ contributions. But what is the nature of this supervenience relation? Focusing on group judgments, we argue that, for a group to be rational, its judgment on a particular proposition cannot generally be a function of the members’ individual judgments on that proposition. Rather, it must be a function of their individual sets of judgments across many propositions. So, knowing what the group members individually think about some proposition does not generally tell us how the group collectively adjudicates that proposition: the supervenience relation must be ‘set-wise’, not ‘proposition-wise’. Our account preserves the individualistic view that group agency is nothing mysterious, but also suggests that a group agent may hold judgments that are not directly continuous with its members’ corresponding individual judgments.

1. Introduction

In this paper, we sketch an account of group agency. We take groups, whether agents or not, to be sets of individuals who are networked with each other in a way that matters to them or others, and that affects their behavior or that of others. The networking may matter because it marks members off in their own perceptions or those of others, or in their capacities or disabilities relative to others; the possibilities are various. Those of a certain religious or ethnic background may form a group on this account, as may those in a particular profession or those with distinctive skills. But those who live at a certain latitude on earth do not form a group, nor do those who are of the same unexceptionable height or hair colour.

What distinguishes group agents from other groups, then? We argue that it is their capacity to mimic the more or less rational way in which individual agents act. Examples of groups constituting agents include committees and commissions, partnerships and companies, expert panels and joint authorships, governments and courts. These groups are not just networked collections of individuals. They are networked collections whose performance parallels that of individual agents. They can take on tasks, commit themselves to goals, enter into contractual relationships,

¹ C. List, Department of Government, London School of Economics, London WC2A 2AE, UK; P. Pettit, University Center for Human Values, Princeton University, Princeton, NJ 08544, USA. Earlier versions of this paper were presented at the NAMICONA Conference on Reductive Explanation, University of Aarhus, May 2005, and at the 2005 Spindel Conference on Social Epistemology, University of Memphis, September 2005. We thank the participants at both occasions for helpful comments and suggestions.

and be held responsible for what they do. They are entities that may have the status of legal persons.

Where does the capacity for group agency spring from? Does it emerge mysteriously, without a clear basis at the level of individuals, as some traditions have suggested (Runciman 1997)? Or does it appear in virtue of how things are organized among individual members? Is it consistent with an underlying individualistic ontology? We explore and defend an individualistic account of group agency here.

If a group is to be a rational agent, under any plausible form of individualism, then it must be constituted in such a way that certain ‘inputs’ by the group members – for example, their actions, judgments or dispositions – give rise to suitable ‘outputs’ at the group level: to outputs that manifest the group’s standing as an agent. The rational agency of the group must ‘supervene’ on the group members’ individual contributions – in analogy to the way in which, on standard accounts, the rational agency of an individual human being supervenes on certain physical processes in this human being’s brain and body. But what exactly is the nature of this supervenience relation? We address this question here.

We argue that the relation required is more complex than might have been expected. Focusing on group judgments in particular, we show that a group’s judgment on a particular proposition cannot generally be a function of the group members’ individual judgments on that proposition. Rather, it must be a function of the group members’ inputs in their entirety. The upshot is that knowing what the group members individually think about some proposition does not generally tell us how the group as a whole adjudicates that proposition. While our account preserves the individualistic view that group agency is nothing mysterious, it also supports the interesting possibility that a group may hold judgments that are not directly continuous with the group members’ corresponding individual judgments.

Our discussion is structured as follows. We suggest general conditions of agency in section 2 and introduce the supervenience account of group agency in section 3. Drawing on the emerging theory of judgment aggregation (e.g. List and Pettit 2002; Pauly and van Hees 2003; Dietrich 2006), we then present some impossibility results in section 4 which show that group agency is not generally consistent with the

requirement of ‘proposition-wise supervenience’. We explore the possibility of group agency under the less restrictive requirement of ‘set-wise supervenience’ in section 5. In section 6 we draw some conclusions. The crucial notions of proposition-wise supervenience and set-wise supervenience will be defined below.

2. Conditions of agency

When does a system, natural or artificial, individual or social, count as an agent? We think that four conditions are individually necessary and at least close to being jointly sufficient. We state the conditions here but do not provide a full-scale defence of them, if only because they reflect a broad consensus in psychology, economics and the philosophy of mind. The conditions are the following:

- First, the system forms representational and goal-seeking states; for example, beliefs and desires, or judgments and plans.
- Second, in forming and revising these representational and goal-seeking states, the system satisfies appropriate conditions of (theoretical) rationality. We will give attention to three such conditions in particular: completeness, consistency and deductive closure, as defined below.²
- Third, the system acts or intervenes in the world on the basis of its representational and goal-seeking states, as conditions of (practical) rationality require; it acts so as to realize its goals, under the guidance of its representations.
- Fourth, the system exhibits these properties not just accidentally or contingently, but robustly – that is, not just in actual conditions, but also in a class of relevant possible conditions.

These conditions should be readily intelligible. Consider a human being, a simple animal, or perhaps a swarm of bees. In each case we can discern a pattern of behavior that invites us to adopt the ‘intentional stance’, as Daniel Dennett (1987) calls it. Once we adopt this stance towards a system, we cannot help but take the conditions above to be fulfilled. We recognize a complexity in the interaction between the system and its environment that leads us to analyze it as a system that more or less rationally

² One might also add certain conditions of truth-tracking.

espouses representations and goals; it acts rationally in accordance with its representations and goals; and it displays these properties more or less robustly, not as a product of fortuitous chance or occasion.

The conditions of agency are formulated in a somewhat abstract way, so as not to engage with unnecessary questions of detail. They say nothing on what internal organization a system must have to count as a rational agent. We may want to stipulate that the system must be wholly present in the spatial boundaries it represents; that it must not be controlled from outer space, for example (Peacocke 1983). We may also want to stipulate that it must generate its responses on the basis of causal connections between successive, evolving states, not on the basis of clever pre-emptive rigging (Block 1980). Both of these qualifications answer to ordinary intuitions (Jackson and Pettit 1990). But beyond those general stipulations, we need to say nothing further on how an agent must be internally constructed. For all we suppose, the architecture of agency may be otherwise unconstrained.

Just as we do not suppose anything specific on this organizational question, so we make no demanding assumptions about how far agents must engage with matters of value. We take it that agents form goals (thereby instantiating states such as plans, desires, preferences, or utilities); agency requires intervention, after all, not just representation. But we can be neutral on the source of those goals. We can preserve our picture of agency, regardless of whether or not we assume that the system's goals are supported by underlying representations to the effect that something is inherently or instrumentally desirable or plan-worthy. Details of our picture may change with changes in our account of these goals, but we need not commit ourselves to any particular account here.

In the following discussion we shall be concerned with how group agents meet one particular necessary condition for agency: that the system robustly satisfy constraints of theoretical rationality, such as the constraint of consistency, in the formation of representational states. More particularly still, we shall be concerned with how group agents can meet this condition with respect to those representational states we describe as 'judgments'. We use the notion of 'judgment' in a broad sense, to include both judgments of fact, bearing on what is to be believed, and judgments of value, bearing on what is to be desired. While 'beliefs' (and 'desires') come in degrees of

strength, ‘judgments’ are categorical. I may believe to this or that degree that *P* but I will judge that *P*, period, or I will not judge that *P*, period. Under what may be a regimentation of common usage, there is no room for holding a judgment more or less strongly. This is not a great restriction, as there is still room for judging that it is more or less probable (or more or less desirable) that *P*.

We focus on judgments because in the case of those group agents we are especially interested in here – such as committees, expert panels, governments, courts, co-authorships – judgments are particularly important representational states. But why focus on judgments rather than plans? Plans are also on-off states, after all, and they also engage constraints like consistency.

We are influenced by the following consideration. Whereas rationality constraints on plans will track corresponding rationality constraints on judgments – judgments of value as to what should be done or brought about – the converse does not hold. We achieve a greater simplicity by focusing on judgments, and we do so without any great loss of generality.


3. The supervenience account of group agency

Under an individualistic ontology, a group’s agency cannot emerge mysteriously without a clear basis at the level of the group members. The ‘outputs’ at the level of the group – here the group’s judgments – must ‘supervene’ on certain ‘inputs’ at the level of the group members. And given the conditions of agency discussed in the last section, the supervenience relationship must guarantee the rationality of the group judgments formed.

We say that one set of facts, *B*, ‘supervenes’ on another set of facts, *A*, if and only if, necessarily, fixing the *A*-facts also fixes the *B*-facts. There is no variation possible in the *B*-domain without a variation in the *A*-domain. An individualistic ontology commits us to the view that a group’s judgments supervene on the contributions of individuals: say, on what the individuals judge and do. More precisely, a group’s judgments supervene on these individual contributions once the group’s ‘constitution’ is put in place. As illustrated in Table 1, a ‘constitution’ is a set of rules, formal or informal, for determining how the inputs of individuals are to be put together to generate group judgments as outputs (see also List 2005). A simple example of a

constitution is the rule that a group judges any given proposition to be true whenever a majority of group members individually judge this proposition to be true.

Table 1: A constitution



In the absence of any constitution, it hardly makes sense to ascribe judgments to a group. The group members' individual contributions are integrated into a group judgment only when an appropriate constitution is explicitly or implicitly in place. Take the people who happen to be in the same subway train at the same time. Clearly there is no formal or informal constitution in place among them, and so it does not make much sense to talk of the group judgments that they hold. By contrast, many groups in public life – such as committees, judiciaries, organizations, companies, expert panels – are organized by appropriate formal or informal rules. And so, at least in principle, they are capable of generating group judgments from individual contributions.

Does the need for a constitution in any plausible supervenience account of group agency compromise the hope for an individualistic ontology? We do not think so. That a constitution is in place among a collection of people merely means that they share certain interpersonally connected dispositions: the dispositions to follow or license certain procedures in the derivation of group judgments from individual contributions. We might think of the constitution, therefore, as yet another individual contribution on the part of the members: a contribution that consists in their possession of the appropriate dispositions. For convenience, however, we shall treat the constitution as a framework within which individual contributions – paradigmatically, judgments and actions – are made and a framework in virtue of which the group-level judgments are formed.

We can now present our main results. If an individualistic account of group agency is to be vindicated, then it must be possible to find a constitution such that the group

judgments generated by it from individual contributions are robustly rational. It must be possible to identify a supervenience relation that is capable of securing this result. We turn now to some results in the recently developed theory of judgment aggregation and explore their significance for this inquiry.

4. Impossibility results: the inconsistency of robust group rationality with ‘proposition-wise’ supervenience

Consider a group of two or more individuals faced with the task of making judgments on some interconnected propositions. In a paradigmatic and much discussed example (Kornhauser and Sager 1986), the group is a multi-member court making judgments on the following propositions:

P: The defendant did action X.

Q: The defendant had a contractual obligation not to do action X.

R: The defendant is liable for a breach of contract.

The propositions are interconnected by the constraint that proposition *R* (the ‘conclusion’) is true if and only if propositions *P* and *Q* (the ‘premises’) are both true: more formally, ‘*R* if and only if (*P* and *Q*)’. More generally, there might be more than two premises; or, in other cases, the disjunction rather than conjunction of the premises might be taken to be necessary and sufficient for the conclusion.

The set of propositions considered by the group – including the logical constraint ‘*R* if and only if (*P* and *Q*)’ – is called the ‘agenda’. Throughout this paper, we assume for simplicity that the agenda is as in the multi-member court example or one of its generalizations, but many other kinds of agendas have been investigated in the literature on judgment aggregation.³ We also assume that, whenever a proposition is included in the agenda, then so is its negation; this enables the group to accept as true either the proposition or its negation or neither.

³ While we here state all formal results just for the agenda of the court example, they can be shown to hold for larger classes of agendas. Proposition 1 holds for all agendas that have a minimal inconsistent subset of three or more propositions; Proposition 2 holds for all so-called ‘minimally connected’ agendas; and Proposition 3 holds for all so-called ‘strongly connected’ agendas. For technical details, see Dietrich and List (2005).

Each group member forms judgments on (some or all of) the propositions in the agenda. We say that an individual's judgments are:

- 'complete' if, for every proposition in the agenda, the individual judges either the proposition or its negation to be true;
- 'consistent' if, for every proposition in the agenda, the individual does not judge the proposition and its negation to be true;
- 'deductively closed' if, whenever the propositions in the agenda judged true by the individual logically entail another proposition included in the agenda, then the individual also judges that other proposition to be true.

Now, given our earlier definition, the group's 'constitution' is a set of rules by which the group members' individual contributions determine the group's judgments on the propositions in the agenda.⁴ We assume in our formal discussion that the group members' contributions are their relevant individual judgments, but in our conclusion below we also consider other possible individual contributions. Moreover, we here assume that the constitution has the 'universal domain': it accepts as admissible input all possible combinations of complete, consistent and deductively closed individual judgments. If that domain is further enlarged so as to include combinations of individual judgments that are not fully rational, our results essentially continue to hold.⁵

What does it mean for the multi-member court in our example to be a group agent? In terms of our necessary condition for agency, the court must form judgments on the propositions in the agenda that satisfy certain rationality conditions.⁶ We can capture this by the following condition, which applies the individual rationality requirements defined above to a group as a whole.

⁴ Formally, a constitution is a function that maps each admissible combination of individual judgments on the propositions in the agenda to corresponding group judgments on these propositions. As noted above, a simple example of a constitution is the rule that the group judges a proposition to be true whenever a majority of the group members judge that proposition to be true.

⁵ Some technical refinements may be needed in this more general case.

⁶ Perhaps additional conditions are required for group agency, but we here consider just a simple necessary condition.

Robust group rationality. The group's judgments (generated through the constitution) are robustly (by which we mean: for all admissible combinations of individual judgments) complete, consistent and deductively closed.

Robust group rationality might seem rather strong: especially completeness and deductive closure seem to be demanding requirements. But notice that completeness and deductive closure are required only for the propositions in the agenda, that is, the propositions on which the group is supposed to make judgments; no such requirements are made for propositions outside the agenda, whose resolution may not be required.

Can the group be constituted in such a way as to meet the condition of robust group rationality? And, if it is, how exactly do the group's judgments supervene on the group members' inputs? A simple and initially plausible thesis about how the group's judgments supervene on these inputs is the majoritarian supervenience thesis.

Majoritarian supervenience. The group judgment on each proposition in the agenda is robustly the majority judgment on that proposition.

But if this is the way in which group judgments supervene on individual inputs, then group agency, in the sense defined above, is not generally possible, as the following result shows.

Proposition 1. For a constitution with universal domain, robust group rationality is inconsistent with majoritarian supervenience.

This result is a slightly generalized version of the much discussed 'discursive dilemma' (e.g. Pettit 2001, ch. 5); for a proof of the present version, see List (2006). To sketch the argument, assume, for a contradiction, that a group of two or more individuals is constituted in such a way that robust group rationality and majoritarian supervenience are both met. By robust group rationality, the group's judgments are complete, consistent and deductively closed for all combinations of individual judgments in the domain of the constitution. In particular, in the special case of a three-member group, consider the individual judgments in Table 2, where the agenda is the one from the court example. This combination of judgments is clearly

admissible under the universal domain assumption; similar examples can be constructed for different group sizes and different agendas of propositions.

Table 2

	P	Q	R	R if and only if (P and Q)
Individual 1	True	True	True	True
Individual 2	True	False	False	True
Individual 3	False	True	False	True
Majority	True	True	False	True

By majoritarian supervenience, the group's judgment on each proposition is the majority judgment on that proposition. But the majority judgments resulting from the individual judgments in Table 2 violate deductive closure: propositions P and Q and the logical constraint ' R if and only if (P and Q)' are each judged to be true by a majority, and these propositions jointly entail proposition R ; yet R is judged to be false by a majority. This contradicts robust group rationality. Notice that this rationality violation occurs despite the fact that the judgments of all group members are individually rational here, in the sense of being complete, consistent and deductively closed.⁷

For a group to be an agent, then, the relation between the group judgments and those of the group members cannot be that of majoritarian supervenience. Could the relation be something similar to majoritarian supervenience? After all, it seems plausible to assume that the group's judgment on a proposition supervenes *in some way* on the group members' judgments on that proposition, albeit not necessarily in a majoritarian way. Consider the following supervenience thesis, which is weaker than majoritarian supervenience.

Uniform proposition-wise supervenience. The group judgment on each proposition in the agenda is robustly a function of the individual judgments on that proposition, where the function depends on more than one individual's judgment and is the same for all propositions.

While the majoritarian supervenience thesis permits only one such function – namely the majoritarian one – the present supervenience thesis permits a large class of functions; it only rules out functions according to which group judgments depend only

⁷ So the inconsistency between robust group rationality and majoritarian supervenience does not depend on any irrationality on the part of the group members.

on the judgments of a single fixed individual. But even if majoritarian supervenience is weakened to uniform proposition-wise supervenience, group agency, in the sense defined above, is not generally possible.

Proposition 2. For a constitution with universal domain, robust group rationality is inconsistent with uniform proposition-wise supervenience.

This result is a strengthened version of an impossibility result in List and Pettit (2002), proved in this strengthened form by Pauly and van Hees (2003). As the proof is more technical than that of Proposition 1 above, we omit it here. But the result shows that the problem illustrated in the sketch proof of Proposition 1 persists even if the group judgment on each proposition is not determined by the majority judgment on that proposition, but by another, more general function of the individual judgments. Again, the result does not depend on any irrationality on the part of the individuals; it is true despite the favorable assumption that individual judgments are rational.

So, for a group to be an agent, the relation between the group judgments and those of the group members cannot be that of uniform proposition-wise supervenience either. Let us relax our supervenience thesis further. Perhaps the problem lies in the ‘uniformity’ of the supervenience relation, that is, the fact that the functional dependence between individual judgments and group judgments is the same for all propositions. Consider the following weakened proposition-wise supervenience thesis.

Proposition-wise supervenience. The group judgment on each proposition in the agenda is robustly a function of the individual judgments on that proposition, where the function depends on more than one individual’s judgment and in addition respects unanimous individual judgments,⁸ but may differ from proposition to proposition.

Proposition-wise supervenience would permit, for example, that on some propositions the group judgment is the majority judgment, while on others it is a different function of the individual judgments. Each such function must only have the specified minimal properties (that is, it must depend on more than one individual’s judgment and respect unanimous individual judgments). But even if we assume proposition-wise

⁸ This means that, whenever the individuals unanimously agree on some proposition, this agreement is respected by the group judgment.

supervenience alone, dropping the ‘majoritarian’ and ‘uniformity’ requirements, we are still faced with an impossibility result.

Proposition 3. For a constitution with universal domain, robust group rationality is inconsistent with proposition-wise supervenience.

Extending an earlier impossibility result by Pauly and van Hees (2003), this result was proved by Dietrich and List (2005); again, we omit the proof. In summary, for a group to be an agent, the relation between the group judgments and those of the group members cannot be that of proposition-wise supervenience. Although this does not refute the supervenience account of group agency, we can already conclude that the supervenience relation cannot be as simple as one might have thought. The group’s judgment on a particular proposition cannot generally be a function of the group members’ individual judgments on that proposition. So if the group is constituted in such a way as to form an agent, the group members’ individual judgments on a proposition are not generally sufficient to determine the group’s judgment on that proposition. The supervenience relation must be more complex.

5. Possibility results: the consistency of robust group rationality with ‘set-wise’ supervenience

The core idea of the supervenience account of group agency is that the rational agency of a group – if indeed the group is an agent in its own right – supervenes on the group members’ individual contributions, here specifically on their individual judgments. In our conclusion, we briefly consider the possibility that the group’s judgments supervene on other, non-judgmental contributions by the group members.

Is group agency ever possible according to this core idea, given that group judgments cannot generally supervene on individual judgments in a proposition-wise way? The following supervenience thesis preserves the core idea of the supervenience account, while giving up the requirement of proposition-wise supervenience.

Set-wise supervenience. The *set* of group judgments on all the propositions in the agenda is robustly a function of the individual *sets* of judgments on (some or all of) these propositions.

We now show that there are possible group constitutions under which a group satisfies both robust group rationality and set-wise supervenience. This finding supports our claim that, at least in principle, group agency is possible under the supervenience account.

Again, consider the multi-member court example. In that example, the group has to make judgments on the propositions P , Q , R , and ' R if and only if (P and Q)' (and their negations). Can it do so in a way that meets both robust group rationality and set-wise supervenience? Consider the following constitution.

The premise-based procedure. The group first makes a group judgment on each premise (here P , Q) by taking a majority vote on that premise (with some constitutional provision for breaking majority ties). The group also accepts the appropriate logical constraint (here ' R if and only if (P and Q)') and then derives its group judgment on the conclusion (here R) from these group judgments on the premises, using that logical constraint.

In our example, the premise-based procedure would require the court first to take separate votes on whether the defendant did action X and on whether he or she had a contractual obligation not to do X , and then to derive its judgment on the defendant's liability from the outcomes of these votes, using the appropriate logical constraint.

Proposition 4. A group using the premise-based procedure as its constitution satisfies both robust group rationality and set-wise supervenience, but not proposition-wise supervenience.

It is easy to see why this possibility result holds (Pettit 2001, ch. 5). First, the premise-based procedure is guaranteed to generate group judgments that are complete, consistent and deductively closed, regardless of the group members' individual judgments: under the premise-based procedure (i) propositions are always decisively adjudicated; (ii) it is impossible for a proposition and its negation to be judged true simultaneously; and (iii) the adherence to the appropriate logical constraint ensures the satisfaction of deductive closure. For example, if the individual judgments are as in the 'problematic' case of Table 2 above, then the premises P and Q are each accepted by a majority vote, the logical constraint ' R if and only if (P and Q)' is

accepted by default, and the conclusion R is accepted by logical implication, an overall rational set of judgments.

Second, under the premise-based procedure, the set of group judgments on the propositions in the agenda is a function of the individual sets of judgments on those propositions: once the individual judgments on all propositions are fixed, the group's judgments are also fixed.

Third, to prove that a group using the premise-based procedure as its constitution violates proposition-wise supervenience, consider proposition R in our example (the conclusion) and notice that the group judgment on R is not determined by the individual judgments on R alone. In particular, there exist two possible situations in which all individuals hold the same judgments on R , and yet the group judgment on R differs between the cases. Compare, for example, the cases of Table 2 (above) and Table 3 (below). The individual judgments on proposition R are the same in these two cases (the column corresponding to R is the same in both cases). Yet, if the group uses the premise-based procedure as its constitution, the group judges proposition R to be true in the case of Table 2 but not in the case of Table 3.

Table 3

	P	Q	R	R if and only if (P and Q)
Individual 1	True	True	True	True
Individual 2	False	False	False	True
Individual 3	False	False	False	True
Majority	False	False	False	True

It is worth noting that the supervenience relation here has not only a set-wise character (as opposed to a proposition-wise one), but also a further property (Pettit 2003). Under the premise-based procedure, the individual judgments on the premises alone are sufficient for determining the group judgments on all the propositions. So the group judgments are non-continuous with the group members' individual judgments in two senses. The individual judgments on the conclusion are not only insufficient for determining the group judgments on the conclusion (a weak discontinuity), but also unnecessary (a strong discontinuity).

The premise-based procedure can be generalized to more than two premises and to other logical constraints (for example, disjunctive rather than conjunctive ones). Moreover, neither the classification of certain propositions as 'premises' and 'conclusions' nor the choice of the logical constraint need to be built into the group's

constitution. A generalization of the premise-based procedure to other agendas of propositions is the following (List 2004, 2006; for informal versions, see Pettit 2001, ch. 5; 2003).

A sequential priority procedure. First, an order of priority among the propositions in the agenda is specified. Earlier propositions are interpreted as ‘prior to’ later ones: they may serve as ‘premises’ in relation to later ones. Second, the group considers the propositions in the given order. For each proposition thus considered, if that proposition is not logically constrained by earlier propositions judged to be true, then the group takes a majority vote on the new proposition; but if the new proposition is logically constrained by those earlier propositions (such as a ‘conclusion’ that is constrained by ‘premises’ judged to be true earlier), then the group derives its judgment on the new proposition from its judgments on those earlier propositions.

It is easy to see that Proposition 4 continues to hold if the premise-based procedure is generalized to a sequential priority procedure. A group using either of these two procedures as its constitution satisfies both robust group rationality and set-wise supervenience, but violates proposition-wise supervenience. Like the premise-based procedure, the sequential priority procedure may give rise to discontinuities between group judgments and corresponding individual judgments.

The premise-based and sequential priority procedures are both constitutions under which all group members contribute to the group judgments in exactly the same way. In particular, if we permute the group members’ contributions, the group judgments are unaffected. For example, if we permute the rows in Tables 1 and 2, the resulting group judgments under the premise-based procedure remain the same in each case. Formally, we say that the supervenience relation between individual judgments and group judgments has a ‘homogeneous supervenience base’.

Homogeneity of the supervenience base. The set of group judgments on the propositions in the agenda is invariant under permutations of the group members’ individual sets of judgments on those propositions.

By contrast, we say that the supervenience relation between individual judgments and group judgments has a ‘heterogeneous supervenience base’ if this condition is violated. Are there any interesting constitutions under which the group judgments

supervene on the group members' individual judgments in a heterogeneous way and where these group judgments are robustly rational? Consider the following constitution (List 2005).

The distributed premise-based procedure. The group is subdivided into multiple subgroups, one for each premise (e.g. one for P and one for Q). Each subgroup 'specializes' on precisely one premise and takes a majority vote on that premise only (e.g. one subgroup specializes and votes on P , another on Q). Now the outcomes of these majority votes are taken as the overall group judgments on the premises. Again, the group also accepts the appropriate logical constraint (e.g. ' R if and only if (P and Q)') and derives its group judgment on the conclusion (e.g. R) from its group judgments on the premises, using that constraint.

In the court example, the distributed premise-based procedure would require subdividing the court into two subgroups, where the members of one subgroup would 'specialize' on the question of whether the defendant did action X and vote only on this first issue, and the members of another subgroup would 'specialize' on the question of whether the defendant had a contractual obligation not to do action X and vote only on this second issue. The members would not necessarily have to form individual judgments on whether the defendant is liable. Rather, the court's overall judgment on the liability issue would be derived at the group level from the judgments reached by the relevant subgroups on the two premises.

In the case of a court, this is an unfamiliar (and perhaps implausible) constitution. However, large committees, and particularly legislatures, are often subdivided into several subcommittees that play exactly the role assigned to subgroups by the distributed premise-based procedure.

Proposition 5. A group using the distributed premise-based procedure as its constitution satisfies both robust group rationality and set-wise supervenience, but not proposition-wise supervenience.

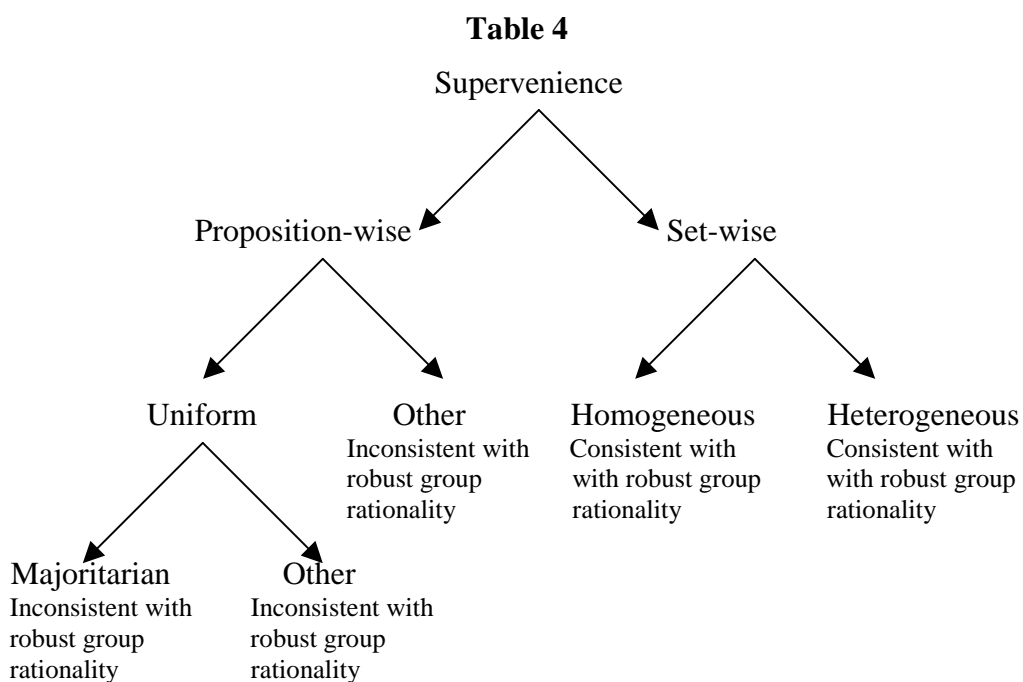
For technical details, see List (2005). The judgments of a group using the distributed premise-based procedure as its constitution are non-continuous with the group members' individual judgments in several senses. First, for each premise, the judgments of a subset of the individuals is sufficient for determining the group

judgment on that premise, whereas the judgments of other individuals are unnecessary; there are different such subsets for different premises. Second, to determine the group judgments on all the propositions in the agenda, each individual needs to contribute only a single judgment on a single proposition, namely on the premise on which that individual ‘specializes’; no contribution on any of the other propositions is necessary. And, third, no individual judgments on the conclusion are necessary for determining the group judgment on the conclusion.

The distributed premise-based procedure in an example of a constitution that allows a group to perform as a unified rational agent based on an internal division of labor.

6. Conclusion

In our formal discussion, we have focused on the question of how the judgments of a group must supervene on those of its members for the group to be rational. A supervenience relation can be proposition-wise or set-wise. Among proposition-wise supervenience relations, we have further distinguished between uniform ones and others, and among uniform propositionwise supervenience relations between majoritarian ones and others. Among set-wise supervenience relations, we have distinguished between cases where the supervenience base is homogeneous and ones where it is heterogeneous. Table 4 summarizes the different supervenience relations we have considered and our formal results.



We have begun with the observation – drawn from the ‘discursive dilemma’ – that a majoritarian supervenience relationship is inconsistent with robust group rationality. But a majoritarian supervenience relation is a highly special one, as it is a special case not only of a proposition-wise supervenience relation, but also of a uniform one. We have seen that even if the restrictions of majoritarianism and uniformity are dropped, proposition-wise supervenience remains inconsistent with robust group rationality.⁹ By contrast, set-wise supervenience is consistent with robust group rationality. The premise-based and distributed premise-based procedures are examples of group constitutions under which group judgments are both robustly rational and set-wise supervenient on individual judgments. Here the supervenience base is homogeneous in the case of the regular premise-based procedure and heterogeneous in the case of the distributed one.

The possibility results are meant to be indicative of how group agency is possible, not exhaustive of the different ways in which it may be achieved. There are a variety of possibilities open, as should be fairly clear. They include informal procedures in which members of the group are invited to think explicitly about the requirements for group consistency, and adjust in the light of these; there is no reason to exclude that possibility (pace McMahan 2005). An example of such a procedure might be the following: individuals take a straw vote on each proposition that comes up, determine whether the straw judgment is inconsistent with existing judgments on other propositions, and then seek to resolve any inconsistency by eliciting a second round of voting on which of the conflicting judgments to revise (List and Pettit 2005; Pettit 2006).

We said above that we would not make any particular assumptions about the internal make-up of agents. The idea was that so long as a system behaves like an agent, it

⁹ It is, of course, possible to identify some special conditions under which (some version of) robust group agency *is* consistent with (some version of) proposition-wise supervenience. With respect to robust group agency, we might relax the robustness requirement of agency, for example by restricting the domain of the constitution. Or we might relax the rationality requirement of agency, for example by relaxing the requirements of completeness or deductive closure. With respect to proposition-wise supervenience, we might permit a trivial proposition-wise supervenience relation whereby the group judgments depend only on a single ‘dictatorial’ individual. Or we might relax the ‘respect for unanimity’ requirement in ‘proposition-wise supervenience’ and permit a trivial proposition-wise supervenience relation whereby the group judgments are held constant across all possible combinations of individual judgments. Finally, we might shrink the agenda of propositions on which collective judgments are to be formed. These possibilities correspond to various escape-routes from the impossibility results on judgment aggregation. See List and Pettit (2002) and List (2005, 2006).

should generally count as an agent. The most that might be required in addition, we suggested, was that the system's responses were not generated from a distant center and that they were not pre-empted by prior rigging. In effect, what we proposed was that function rather than structure is what matters for agency.

We have illustrated in the later sections of the paper ways in which a group's structure or constitution may vary while group agency is preserved: in particular, while the rationality of the group agent's judgments is preserved. But the question that naturally arises, in conclusion, is whether we have pointed at the further reaches of possibility in this domain. Does our approach make room for all the possible ways in which individuals might cooperate with one another to constitute a group agent?

In rounding off this discussion we have to admit that we may have been too conservative in one respect. Especially in our examples of possible group agents, we have implicitly assumed that the individuals who constitute a group agent do so in a knowing and willing manner. They do so, if they do it in full-dress form, on the basis of certain 'joint intentions' (Tuomela 1995; Bratman 1999; Gilbert 2001). Each member intends that together members sustain the group agent in operation; members will at least acquiesce in the more or less salient fact that how they act together secures that result. And, regimenting their attitudes in full dress, each member intends to do his or her bit; believes that others will do their bit; intends to do his or her bit because of this belief; where all of this is above board, as a matter of shared awareness (Pettit and Schweikard 2006).

Might individuals ever constitute a group agent without anything, however implicit, approximating this condition? The question can be sharpened with an example from non-human animals. It is often said that a swarm of insects can behave as if it were a single, organized agent, even though each individual insect presumably responds in a more or less rote way to chemical signals from its neighbors or environment (Seeley 1989). The swarm, we may suppose, behaves like a proper agent, the individual bees with the inflexibility of automatons, and so without any awareness of the swarm-level behavior. Can we imagine human beings constituting a group agent on a similar basis: on a basis that does not require any one of them to have the conception of what they as a group are doing?

Some of Tolstoy's discussions in *War and Peace* suggest that he thought of populations having this sort of emergent agency, without individuals really understanding what was going on. But we remain skeptical about the possibility. The requirements for such emergent agency look, on the face of it, to be implausibly strong. The individuals who contribute to the group in action will presumably do so, at least in some part, by acting in their own right. But if they act without a conception of their contribution to the group, then their reasons for action must be unrelated to the group's performance. And in that case it is not easy to see what sort of organization, what sort of unrecognized constitution, could guarantee that group agency would be secured. How could any constitution ensure that no matter what people's personal reasons for acting, they will always act as is required for the group as a whole to be robustly rational (Pettit 1993, ch. 3)?

It is hard to see how a constitution could do this, unless the whole enterprise was directed centrally and members deferred to the director. Take Ned Block's China-body system (Block 1980, pp. 276-277). In this imaginary scenario each of the billion members of the Chinese population takes charge of a particular task in the Turing-machine replication of someone's mental life. Without individuals understanding what they are doing, their electronic connections with one another and with the artificial body through which they act ensure that the body manifests agency. Does this mean that the members of the population constitute a group agent? Perhaps, but the presence of central direction, and the widespread deference to the director, would mean that we have a special sort of joint intention here: an intention on the part of each that they together follow what the director enjoins.

If standard social and economic theory is to be believed, then the individually rational inputs of individuals can generate, as by an invisible hand, a pattern of collectively rational results. The question is whether a group agent might emerge in the same way. And it is not clear to us how, empirically, it could. Collectively rational results – say, a pattern of competitive pricing – are stable across many variations in context. But the outputs that would have to be generated for the emergence of a group agent will have to be tailored to different circumstances of action. No existing theory makes sense of how this could happen.

Donald Davidson once said that the secret in exploring a philosophical thesis is to maintain the excitement while increasing the intelligibility. We have explored the thesis that rational group agency supervenes, but not in a straightforward way, on the contributions of individual members. The most exciting version of that thesis is certainly the doctrine of emergent agency that Tolstoy supports. But at this margin of excitement, alas, the intelligibility runs out. We have to settle for less. Such is life.

References

- Block, N. (1980). "Troubles with Functionalism?" In *Readings in Philosophy of Psychology, Vol 1*. London: Methuen.
- Bratman, M. (1999). *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- Dennett, D. (1987). *The Intentional Stance*. Cambridge/MA: MIT Press.
- Dietrich, F. (2006) "Judgment aggregation: (im)possibility theorems." *Journal of Economic Theory* 126(1): 286-298.
- Dietrich, F. and C. List (2005). "Arrow's theorem in judgment aggregation." Working paper, University of Konstanz.
- Gilbert, M. (2001). "Collective Preferences, Obligations, and Rational Choice." *Economics and Philosophy* 17: 109-120.
- Jackson, F. and P. Pettit (1990). "In Defence of Folk Psychology." *Philosophical Studies* 57: 7-30; reprinted in F. Jackson, P. Pettit and M. Smith, *Mind, Morality and Explanation*. Oxford, Oxford University Press (2004).
- Kornhauser, L. A. and L. G. Sager (1986). "Unpacking the Court." *Yale Law Journal* 96(1): 82-117.
- List, C. (2004). "A Model of Path-Dependence in Decisions over Multiple Propositions." *American Political Science Review* 98: 495-513.
- List, C. (2005). "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective." *Episteme: A Journal of Social Epistemology* 2(1): 25-38.

- List, C. (2006). "The Discursive Dilemma and Public Reason." *Ethics*, forthcoming.
- List, C. and P. Pettit (2002). "Aggregating Sets of Judgments: An Impossibility Result." *Economics and Philosophy* 18 (1): 89-110.
- List, C. and P. Pettit (2005). "On the Many as One." *Philosophy and Public Affairs* 33(4): 377-390.
- McMahon, C. (2005). "Pettit on Collectivizing Reason." *Social Theory and Practice* 31(3): 431-450.
- Pauly, M. and M. van Hees (2003). "Logical constraints on judgment aggregation." *Journal of Philosophical Logic*, forthcoming.
- Peacocke, C. (1983). *Sense and Content*. Oxford: Oxford University Press.
- Pettit, P. (1993). *The Common Mind*. New York: Oxford University Press.
- Pettit, P. (2001). *A Theory of Freedom: From the Psychology to the Politics of Agency*. Cambridge and New York: Polity and Oxford University Press.
- Pettit, P. (2003). "Groups with Minds of their Own." In F. Schmitt (ed.), *Socializing Metaphysics*. New York: Rowan and Littlefield: 167-193.
- Pettit, P. (2006). "Participation, Deliberation and We-thinking." In D. O'Neill, M. Shanley and I. Young (eds.), *The Illusion of Consent: Essays in Honor of Carole Pateman*. Philadelphia: Pennsylvania State University Press.
- Pettit, P. and D. Schweikard. (2006). "Joint Action and Group Agency." *Philosophy of the Social Sciences* 36.
- Runciman, D. (1997). *Pluralism and the Personality of the State*. Cambridge: Cambridge University Press.
- Seeley, T. D. (1989). "The hone bee colony as a superorganism." *American Scientist* 77: 546-553.
- Tuomela, R. (1995). *The Importance of Us*. Stanford/CA: Stanford University Press.