Routledge
Taylor & Francis Group

INQU Inquiry
An Interdisciplinary Journal of Philosophy

INQUIRY
An Interdisciplinary Journal of Philosophy
vol. 66, no. 4, May 2023
Routledge

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/sinq20

Five elements of group agency

Philip Pettit

To cite this article: Philip Pettit (2023): Five elements of group agency, Inquiry, DOI: 10.1080/0020174X.2023.2208391

To link to this article: https://doi.org/10.1080/0020174X.2023.2208391

INQUIRY
https://doi.org/10.1080/0020174X.2023.2208391

Routledge
Taylor & Francis Group

Five elements of group agency

Philip Pettit[a,b]

[a]Princeton University, Princeton, USA; [b]Australian National University, Canberra, Australia

ABSTRACT
Group agency requires a number of people to combine in pursuit of shared goals across varied scenarios. Thus, a group or corporate agent must be organized (1) to act flexibly as its goals require, (2) with the intentional, if not always voluntary, acquiescence of members in the guidance of (3) an authorized spokesperson or (4) a constructed voice, thereby (5) becoming capable of making and honoring commitments.

ARTICLE HISTORY Received 24 April 2023; Accepted 25 April 2023

KEYWORDS Group; agency; goal; spokesperson; commitment

When a group agent acts to achieve some end, then there is bound to be a sense in which the members act jointly for the realization of that end. But if a group is to act in the manner of an agent, it is not enough that the membership perform a collective or joint action together. In this paper, I try to identify some crucial conditions that the group must also satisfy.

Group agency, as I use the phrase, is the agency exercised by a number of human beings when they combine to pursue shared goals across varied scenarios. Elsewhere I have referred to it as incorporated agency but that term may suggest a focus only on corporations: commercial businesses or firms. Corporations certainly display group agency, as they involve a number of people, usually operating via different subgroups, who combine across different situations in pursuit of corporate goals such as the provision of certain services or the production of certain goods and ultimately the generation of suitable profits. But there are a variety of other group agents in our societies as well: churches, political parties, trade unions, social clubs, and the like; indeed, the states

CONTACT Philip Pettit philip.pettit@anu.edu.au University Center for Human Values, 305 Marx Hall, Princeton University, Princeton, NJ 08544, USA

that rule over our societies and represent them in the international world operate as group agents too (Pettit 2023).

There are five elements of group agency that I would like to distinguish in this short piece. They require respectively: purposive flexibility, individual acquiescence, authorizing a voice, constructing a voice, and having a capacity for commitment.[1] The paper is in five sections, corresponding to those elements, and is summarized in a short conclusion.

1. Purposive flexibility

By all accounts an agent is a system that acts to bring about certain goals or purposes, whether these be fixed or flexible. The agent will typically generate such results because of having been selected or designed to do so or because of having been selected or designed to follow procedures that identify the goals to realize. But agency requires more than the pursuit of purposes. Otherwise the sunflower plant that tracks the sun – that reliably moves so as to maintain its orientation towards the sun – would count, implausibly, as an agent.

The most obvious requirement of agency over and beyond the pursuit of certain goals is that this pursuit should be maintained over variations in the circumstances of the system: variations that are more significant, intuitively, than those in the angle of the sun that the sunflower has to cope with. A system may count as an agent and be so simple that it has only a single goal. But still, it must be able to pursue that goal, acting for its realization, across different scenarios; it must be flexible in the behavior it adopts to realize the goal, adjusting as appropriate for its purposes in different circumstances.

Take the simple robot that is constructed to pursue the goal of putting certain objects on a flat surface in an upright position. Such a robot will have to be equipped with some perceptual apparatus for determining whether any glass or cup or bottle on the surface – say, a table top – is on its side or upright; think of this as a set of receptors that continually scan the objects on the surface. And then, presented with a bottle on its side, it will have to be organized across variations in the size and shape of the object, and its distance and direction, to adjust behaviorally so as, things going well, to put it upright: this, presumably, with the help of wheels for moving, mechanisms for grasping, levers for lifting, and so on (List and Pettit 2011).

[1]The approach taken is in line with my work with Christian List in List and Pettit (2011). For an alternative, broadly congenial approach to group agency, see Tollefsen (2015).

In order for an agent like the robot to adjust its behavior in different scenarios so as to realize one or another goal, it must be able in some sense to register the particularities of each scenario. It must change in response to perceptual inputs, and its changed state must then serve to shape the behavior – and to shape the behavior as it evolves – so as to achieve the goal. This is to say, in other terms, that the agent must not just be disposed to pursue certain goals or purposes; it must be disposed to pursue them according to representations of the environment that it forms. The representations that guide the agent in this way may be of any of a variety of kinds, for all that we need assume here.

The lesson, in familiar Humean terms, is that an agent is a system that reliably acts to fulfill its desires according to beliefs that it reliably forms; reliability in the first dimension constitutes epistemic rationality, we might say, reliability in the second practical rationality. Or it is at least system that is reliable in these ways when circumstances are normal by independent criteria. The robot in our example may be misled under certain lighting about the position of an object, failing in an epistemic way, and it may not manage to put an object at the edge of the table upright – it may knock it to the floor – failing in a practical manner.

Like any agent, a group agent will have to have a goal or, more likely, a set of goals that it reliably pursues across different scenarios according to representations of the particularities of each scenario. It will have to be organized so as to adjust behaviorally as those particularities require, taking steps that reliably, if not inevitably, realize a relevant goal. It will have to display purposive flexibility, in other words, having the capacity to track its purposes in accordance with its representations.

The purposive flexibility required of a group agency will be absent in a group of people who act jointly for a given goal but only in a more or less determinate environment. Think of the people on the beach for whom it is manifest that they recognize that a child is in danger in the water; that they each desire to save it from drowning; that no one is able or willing to try to rescue it on their own; that there is a salient plan – forming a chain into the water – whereby they can together save the child; and that anyone acting on that plan will be joined by sufficient others. Those people will be led to act so as together to save the child, thereby displaying joint or shared agency (Bratman 2014). But such shared agency is not yet group agency (Pettit and Schweikard 2006).[2]

[2]There are many accounts apart from Bratman's of what shared agency involves. See, for example (Gilbert 2015), (Tuomela 2007), Searle (2010) and for an overview of the issues between such writers (Ludwig 2016).

The beach group do perform an action together, each acting so that they as a group succeed in bringing about a certain effect: the rescue of the child. But they are not organized so that they would more or less reliably act for that or any other purpose across what intuitively counts as a significant variety of situations. They are not organized to be purposively flexible.

We might plausibly ask what the group would have done if there had been a life raft on the beach or, more radically, what it would have done if they were hiking and the child was in danger on a crumbling cliff edge. But there is nothing about the group as such – nothing about how it is organized – that would enable us to answer such a question or even to believe that it has a determinate answer. We could only rely on working out what sort of plan, if any, such helpful individuals would be likely to identify in such a variant scenario.

This is to say that unlike a group agent like a corporation or church or club, the beach group lacks purposive flexibility. The members act together episodically for the achievement of a particular goal – saving the child – working out an ad hoc plan for doing so. But they do not have the sort of higher-level plan – a plan, as it were, for making plans – that would enable them as a group to work out episodic plans for any in a range of scenarios.[3] They do not have purposes that transcend the beach scenario and even if they did, they do not have any established procedures for working out how to pursue those purposes under different circumstances. They are not organized for action in the manner of an agent.

There is a sharp contrast in this regard with the corporation or church or club. Any such body will be established to act for certain goals or purposes across a range of variable circumstances – these commercial, ecclesiastical or social goals may themselves be fixed or fluid – and will be organized under certain procedures to determine how best to advance the purposes, and in what order, under different circumstances. We will not only ascribe attitudes like beliefs, desires and intentions to the body in a given scenario, as we can with the beach group. We can also believe that there is a fact of the matter, albeit a fact that is indeterminate in certain ways, about how it would adjust such attitudes and about how it would therefore act in counterfactual variations on the scenario: this, at any rate, provided the variations are within the range of situations with which the group is organized to cope.

[3]They do not have what Christian List and I (2011, 60) characterize as an organizational structure; this corresponds, broadly, to what Peter French (1979) describes as a body's internal decision structure.

This is to say, in decision-theory terms, that a group agent like a corporation has a kinematics: a basis on which it is disposed to update its attitudes in response to variations of circumstances, doing so robustly over changes in or of its personnel. It may not survive long enough to update its attitudes in time but it must be able to do so across a range of modal possibilities. This is to say, in at least one sense, that it must have a set of attitudes and associated habits of inference that constitute an interconnected, relatively holistic web.

2. Individual acquiescence

The network of cells that constitute any biological agent is an example of an organic, not a group, agent. The cells are non-agential entities but provide the components out of which the network is constructed. In virtue of the interaction of those cells with the environment, with one another and with its movements, the network is perceptually sensitive and behaviorally responsive, on a pattern that allows us to ascribe various attitudes to it in this or that scenario. And in virtue of the interaction of its cells, the network is organized within itself so that it reliably updates those attitudes in response to new inputs.

Organic agents may also be constructed out of components that are themselves agents and yet resemble the biological agent in a striking way. Take the colony of ants or the hive of bees (Hölldobler and Wilson 2009). These groups act for discernible goals in the manner of a single biological agent, despite the fact that we think of the component members as themselves agents in their own right. But the things they do as individual agents are programed on a more or less mechanical basis – on a basis over which they have no individual discretion, as it were – in the manner of the cells in a single organism. And those behaviors have the aggregate effect of maintaining the colony or hive in existence over all sorts of changes in surrounding circumstances: for example, in the sorts of sites available for nesting, in the kinds and distances of food sources, and in the dangers that they face from enemies or from nature. Or they have that effect, at any rate, under what we take to be normal circumstances (Seeley 2001).

Might human beings come to constitute a group agent in the ways in which these insects do so? Or might they constitute such an agent on certain occasions, as with a flock of birds foraging for food across a variety of environments (Couzin and Levin 2015)? Assume that the flock behavior is generated by more or less automatic dispositions: for

example, by two interacting disposition in each bird, one to move towards any food source detected, the other to keep within a certain distance of neighboring birds. Might human beings ever constitute group agents by virtue of that sort of sub-personal alignment?

My assumption is that the answer is negative. We human beings may be susceptible to all sorts of subconscious priming and prompting but for anything we do that might contribute to the achievement of a group purpose, we think of it as something for which we can in principle be held responsible. We think of it as something for which we can individually be blamed or, if blame is inappropriate, as something that we can justify or at least excuse. We think of it, in short, as our individual action, whether an action adopted in full knowledge and with full intent, or adopted in ignorance or under pressure. The main point is that we do not think of it as a reflex or automatism over which, necessarily, we have no control (Pettit 1993, Ch 3).[4]

Assume, then, that we are not mobilized into acting for group agents by an automatic, unrecognized mechanism like that which may operate in the flocking birds. Assume, more generally, that we are not mustered into suitable behavior by a device that operates behind our backs. How in that case do we human beings get to be able to constitute group agents?

The only possible answer is, by acquiescing in playing the part required of us by the organization of the group and by accepting the role of other agents in playing their parts too: by going along actively and passively with the requirements that the agency of the group imposes. As the member of a corporation, whether a shareholder, a director, a manager or a worker, I must acquiesce in what I am required personally to do at the AGM, in board meetings, in organizing initiatives or on the shop floor. And, subject to complaints or proposals I may be allowed to raise, I must accept the right of others to behave in their roles as I behave in mine. However well or badly I understand the way the organization works, I must assume that there is a more or less reliable set of procedures in place and must accept that I am contributing in some measure to the overall pattern of group activity.

Group agents operate by the intentional acquiescence of all its members, be they founders or latecomers, in going along with what it requires of them and their fellows. This is scarcely surprising news,

[4]This may be to reject the idea thaat a social fact, in Durkheim's (1938, 10) account, 'is to be recognized by the power of external coercion which it exercises or is capable of exercising over individuals'; he often suggests that the external coercion involved operates subconsciously and in a way that undermines our intentional control of what we do.

being a matter of common observation. But it marks a sharp contrast with the group agents that insects and other animals form in the sorts of examples mentioned.

The individual acquiescence required of members of a group agent may sound faintly reminiscent of contractual theory. But that would be misleading, at least in one respect. For the acquiescence of someone in the operation of a group organized for agency may not be voluntary, although it must be intentional.

Let an option, O, be a possibility taken under one or another description, where it is up to me as an agent to realize that possibility or any of the alternative possibilities that go to constitute the choice before me. I will take an option intentionally, in the strongest sense of the idea, insofar as I adopt it knowingly and deliberately: I act with the intention of realizing the possibility under that description.

But the option O, as I foresee, may have a consequence C. And foreseeing that consequence, I may also be said intentionally to bring about, not just the package, O + C – this is unproblematic – but the consequence, C, itself. The sense of intentionality that attaches to such a consequence is secondary, however, and weaker than the other. This is because I may be said to bring about C intentionally, even when I have no wish or intention to bring it about as such: even when I satisfy the knowingly-condition but not the deliberately-condition (Scanlon 2008, 8). Suppose I knowingly and deliberately go climbing in dangerous mountains: suppose I go climbing intentionally in the strong sense. I may foresee that this will put my health at risk but, while regretting this, still go ahead with the venture. In that case I may be said intentionally to put my health at risk, where I do so knowingly but not deliberately.

It is hard to imagine human beings who act for a group agent but do not do so intentionally in at least the weaker sense. They may follow someone's instructions for the sake of pleasing that person or earning a wage or whatever, where acting in that way means that they play a part in sustaining a group agent and advancing its ends. But we would hardly regard them as a member of the group – we would hardly take the group to be an agent that they help with others to constitute – unless, at the least, they foresaw that acting as instructed would contribute to the group's performance.

But while the members of a group agent must play their parts intentionally in the constitution and performance of the group, they need not do so voluntarily. As intentionality can be taken in two senses, so the same is true of voluntariness.

I may be said to do something, X, voluntarily insofar there is apparently at least one alternative, Y, that is comparatively acceptable: it is acceptable to the point where its presence means that I can be censured or commended – I can be held responsible – for doing X; after all, I could have done Y instead. Taking an option is voluntary in a first, stronger sense just in case this condition is satisfied. The action, as we might say, is discretionary; it is up to me whether to take it or not.

To introduce the second, weaker sense of voluntariness, imagine that while I do not have discretion over whether to do something, because all the apparent alternatives are comparatively unacceptable, still I do it with relish: I do not do it just for want of an acceptable alternative. Imagine that I am coerced, on pain of being shot, to torture someone but I take pleasure in doing so. Do I torture the person voluntarily? Well, not in the sense that I have discretion about whether or not to do it. But I do it voluntarily in the distinct, weaker sense of at least doing it willingly.

While the members of a group agent must intentionally acquiesce in sustaining the group, at least in the weaker sense of intentionality, it should be clear that they may not do so voluntarily in either of these senses. They may assume their role in a context where the alternative is comparatively unacceptable: this, in the way workers may sign up for employment in a corporation when the alternative is starvation. In that case, their participation is bound to be non-discretionary and involuntary in the corresponding sense. This non-discretionary activity, however, may or may not be involuntary in the sense of being unwilling. If the employer offers generous, non-exploitative wages and conditions, then they will probably sign up quite willingly; if the employer makes an exploitative offer, they are likely to sign up unwillingly.

For all that group agency requires, as indeed this example shows, the individual acquiescence of members may not be discretionary or even willing; it may be involuntary in a double sense. But still, the contrast with the case of group agents among insects and other animals remains. Group agents appear among human beings only on the basis of the intentional participation of members, even though the participation elicited may be non-discretionary and unwilling.[5]

[5]This observation has implications for the degree to which, and the ways in which, the different members of a group agent can be held responsible for the actions of the group; see (Pettit 2007; List and Pettit 2011, Ch 7).

3. Authorizing a voice

Assume, on this basis, that group agency requires members to acquiesce in going along with the group, playing their parts and accepting the right of others to play their parts in turn. That leaves a question, however, as to what it is they are required to acquiesce in and we should turn to that issue now.

Whatever the object of their acquiescence, it must identify the parts that different members have to play and, in doing this, it must make it possible for the body as a whole to merit the status of an intentional agent, with the attitudes and the kinematics that agency requires. What sort of object might be able to meet that design specification, playing the role required?

For a cue, we may turn to Thomas Hobbes (1994, Ch 16). He argues on the basis of older, medieval ideas that a straightforward way in which a group agent might form and operate, be it a company of merchants or a commonwealth – these are his examples – is, first, to establish a single spokesperson; second, to let that individual determine and, when appropriate, update the purposes and representations with which the group is to operate; and, third, to distribute to members, individually or in collectivities, the tasks that will enable that group to act for those purposes according to those representations.

The spokesperson envisaged by Hobbes would provide the group with a voice that, speaking for the whole, spells out the general purposes of the body and the general representations by which it should be guided, and then establishes protocols under which it should adjust to particular circumstances. The protocols for adjustment to any scenario would identify the initiatives the group agent should take in pursuit of its goals there, the representations by which it should be guided in that pursuit, and the particular members of the group who should act on its behalf in taking those initiatives. The spokesperson might adopt a hands-on version of control, playing a direct part in applying the protocols, or might opt for arm's-length control: it might let standard routines govern adjustments, oversee what is done from a background position, and remain on standby, ready to intervene if problems arise.

On this picture, the words of the spokesperson would serve a double role. First, they would spell out, at least at a high level, the purposes and representations that the group serve; the spokesperson would give expression to the general attitudes that ought to be ascribed to the

group: the desires and intentions and beliefs that it generally endorses.[6] Second, assuming a hands-on case, they would instruct members – individually acquiescent members – about how they should act in any scenario when acting in the name of the group; or, an arm's-length case, they would instruct members in such a scenario about what routines they should follow in determining how to act.

The voice of the spokesperson, under this simple model of how a group agent might form, is authoritative in character, binding the members insofar as they do not opt for leaving the group or trying to initiate a revolt within. When the spokesperson speaks in this authoritative way for what the group seeks and thinks, the words uttered are not reliable, assuming they are reliable, in virtue of corresponding accurately to the independently established attitudes of the group. They are reliable, rather, because the fact that the spokesperson ascribes those attitudes in speaking for the group makes them into its attitudes. The spokesperson's say-so determines what the group seeks and thinks: it dictates those attitudes in the fashion of an authority rather than describing them in the manner of a reporter.

There is no mystery about why the spokesperson's voice might have this authority in speaking for the group. The individual acquiescence of members in that voice involves the expression of an intention to behave as the group requires them to behave when acting in its name. How the group requires members to behave is determined by the voice that speaks for the group: they should behave in a manner that vindicates the ascription of group attitudes implicit in the spokesperson's words. By acquiescing in the role of the spokesperson in directing them in that manner, the members authorize the spokesperson's voice: they give it the status of being their collective voice.

The simple dictator model of a group agent is not of interest in itself, since few bodies operate under the rule of a single, authorized spokesperson. It is also misleading in the sense the entity that forms under its aegis might be cast, not as a group agent that organizes around a dictator, but as an individual agent, the dictator, who achieves an empowered status by being able to get others to acquiesce in pursuing his or her desires according to his or her representations.

[6]If the spokesperson operates in a hands-on mode, then the words may even serve to express the more specific attitudes on which it acts in this or that particular scenario. Otherwise those more specific attitudes will be there to be read off from the group's behavior, and from the spokesperson's profession of its general attitudes.

But the simple model is nonetheless instructive, for it points us towards a more general possibility. This is that a group agent might get to form as a result of the members establishing a voice in which they can individually acquiesce. This need not be the natural voice of a single, dictatorial individual but an artificial voice that they combine to create. It will do the job of the spokesperson insofar as it is a voice that they can each follow in their individual efforts in the name of the group and so a voice that can serve to determine and update the attitudes maintained by them as a group.

4. Constructing a voice

Hobbes himself thought that it is perfectly possible for some individuals to organize themselves as an agent by creating an artificial voice in whose directive role they can all acquiesce. He maintained in line with received thinking, and in line with the later thinking of Locke and Rousseau, that individuals might organize as a group agent by each consenting to be guided by the majoritarian voice of the membership as a whole or, though we can put this aside, by the majoritarian voice of a proper subset of members (Hobbes 1994, Ch 16; see too Pettit 2008).

Hobbes assumed that the members could vote on every issue that the group might confront, from general questions about the aims it should pursue and the assumptions it should make, to more specific questions about how to act on those aims and assumptions in this or that scenario. And he argued that they could individually acquiesce in being guided in their group efforts by the majority vote, or at least by the majority vote amended to deal with the possibility of ties. That vote would establish a voice, so he assumed, that they could follow in the manner of the voice of a single spokesperson: it would be the voice of the membership, being formed in equal responsiveness to the views of each.

Although Hobbes does not comment a lot on such details, his majoritarian proposal does not strictly require that an assembly of all members should gather to decide on every issue before the group. For example, the proposal might be developed in a model where the assembly votes only on general matters: the framework or constitution of the group; consigns more particular issues to the majority decisions of subgroups it establishes; and retains only an arm's-length form of background oversight and standby control. Thus, we need not dismiss the proposal on the grounds that the assembly requirement that seems to be built into it is simply infeasible.

No matter how it is developed, however, the majoritarian way of identifying an artificial voice for a group to follow is infeasible in another, deeper way. While the majoritarian voice is appealing insofar as it is responsive to individual members, it is unavailable for the role of guiding a group collectively. This is because it is liable to force the group to lock into irrational attitudes that no agent could be guided by: to adopt those attitudes in a way that makes it insensitive to the sorts of problems irrationality generates, and impervious to the criticism that such problems are likely to prompt. This is the lesson of the discursive dilemma (Pettit 2001a).[7]

The dilemma consists in the fact that if a group of individuals tries to operate under the direction of a majority voice, they are liable to have to face a hard choice between being individually responsive and being collectively rational. They may let the voice of the group form in majoritarian responsiveness to the attitudes of members. Or they may let the voice of the group assume a rational profile of a kind that members can follow without problems in acting for the group. But they cannot be sure of being able to do both at once.

The reason why a group may be forced by majority voting into endorsing irrational attitudes can be illustrated with the help of a group of three individuals, A, B and C, voting in a majoritarian way on three connected issues: whether p, whether q, and whether p&q. Only A and B may vote 'yes' to 'p', only B and C may vote 'yes' to 'q'. And so, despite the fact that the group gives majority support to the judgment that p and the judgment that q, they will find themselves voting as a majority against the conjunction: in effect, voting for the judgment that not-p&q.

The discursive dilemma shows that there is liable to be a problem for any group that hopes to determine the voice by which it is guided on a wholly majoritarian basis. But other impossibility results have shown that a similar problem is liable to arise with roughly any procedure for identifying such a voice that, like majority voting, is bottom-up in the following sense: it takes the views of individuals on any issue before the group – whether p, whether q, and the like – and uses them to fix the view of the group as a whole on that issue.[8]

[7]The discursive dilemma is a generalization of the discursive paradox that can arise in judicial hearings; see (Kornhauser and Sager 1993). On the connection, see (Pettit 2001b).

[8]An early theorem is in List and Pettit (2002); it also lays out the other conditions that a bottom-up procedure needs to satisfy if the theorem is to hold. For an informal review of results in the area, see List and Pettit (2011) and for a symposium of relevant, technical papers see (List and Polak 2010). The relation between these results and Arrow's impossibility theorem is discussed in List and Pettit (2004) and Dietrich and List (2007).

But while any bottom-up procedure for identifying a suitable group voice is liable to fail in this way, there is an easy fix for the problem. This is to construct the voice in a way that combines a bottom-up and a top-down element. Insofar as the procedure is bottom-up it will have the appeal of being relatively responsive to the views of individual members on particular issues. Insofar as it allows for top-down intervention the procedure can be adjusted so that this responsiveness does not produce collective irrationality.

The easiest fix involves a straw-vote procedure. It would require the members of the group to look first at what bottom-up voting would generate in response to each issue; to consider in a top-down way whether accepting that vote would generate an inconsistency; and, if it does generate inconsistency, to resort to discussion in order to decide on which vote to reject: one of the prior votes, or the vote just taken.

The procedure can be illustrated with the example of our three individuals, A, B and C, voting on whether p, whether q and whether p&q. The members may ensure that the voice they construct is consistent or rational, and so capable of guiding them in acting for the group, if they follow these steps:

(a) Vote bottom-up on each issue that arises.
(b) Check top-down on whether that vote is consistent with prior votes
(c) If it is consistent with prior votes, adopt it, at least provisionally (see e).
(d) If it is inconsistent, identify the minimal set of inconsistent votes.
(e) Resolve the inconsistency by rejecting one or more of those votes.

Applied to our example, the straw-vote procedure would allow A, B and C to endorse the initial vote on whether p and the second vote on whether q. But it would force them to reconsider their position at the point where they vote that not-p&q. That vote is inconsistent with the prior votes, and the inconsistency would force the members to consider how to resolve the problem. They might do this by rejecting the initial vote, the second vote or the vote on the conjunction. Or of course they might fail to resolve it, with no one willing to revise their vote; in that case, they would fail as a purported group and might be forced to dissolve.

While the straw-vote procedure illustrates the strategy of mixing bottom-up and top-down elements, it should be clear that there are many ways in which the strategy might be implemented. The group

might establish a separate committee of members to do the top-down check on whether its majority voting at any point leads to an inconsistency and if it does, to call for reconsideration. Or it might delegate the reconsideration itself to that committee, enabling it to recommend or even dictate the change to make. And of course, it might take a very different tack by getting a committee to work out the consistent combinations possible, in a top-down initiative, and then letting the membership as a whole decide between them; in the example given, the different combinations possible would be: not-p, q, not-p&q; p, not-q, not-p&q; not-p, not-q, not-p&q; and p, q, p&q. These are just a few examples of the myriad ways in which the strategy of mixing bottom-up and top-down devices might be realized.

There is an interesting contrast between a group's constitution under the dictator and majoritarian models sponsored by Hobbes and under any model in which the group's decisions are made in a way that introduces a top-down element. The voice of the dictator or majority is a voice that is established independently of the existence or operation of the group. The dictator's voice is independently determined by that person's will and appropriated by the group. The majority voice is independently determined as an algorithmic function of the views of individuals and is appropriated in the same way. In each case, the voice adopted by the group is ready-made – it is available off the rack – and the group takes it over more or less passively as the voice to authorize in its guidance.

The voice that the group adopts under a procedure with a top-down element is not taken from elsewhere but is constructed, directly or indirectly, by the group itself. Thus, in our simple example, A, B and C resolve the problem of what to think about the three issues they confront by taking the inconsistent package of majority proposals as material to work on and then constructing the voice by which they are to be guided out of that material.

The voice they endorse will have to differ from the majoritarian package in at least one element, be that p or q or not-p&q; it has to endorse a proposition that a majority rejects. This means that it will have to be a relatively autonomous voice, and not a reflection, issue by issue, of what members think or vote. And as that voice will have to be autonomous in this sense, so it will have to be innovative, representing a relatively free choice on the part of the members. This shows up in our simple example, since the group may choose to reject the majoritarian package in favor of any of a number of the consistent combinations

possible. The voice of the group will be constructed in the double sense of not being systematically determined issue-by-issue by the views of its members and of representing an innovative choice on their collective part.[9]

Assume, then, that a group agent has to construct the voice by which it is to be guided – the voice in which members individually acquiesce – in this active manner. And grant, on the basis of previous considerations, that this voice will authoritatively dictate the attitudes that the group acts on: the mind that it displays. It follows that any group agent will have to make up its mind about things as it goes along: it cannot take its mind as ready-made in the manner of a dictatorial or indeed a majoritarian mind. It may make up its mind in one way or another and whatever mind it forms, that will be a mind of its own (Pettit 2003).

In conclusion, an observation that applies whether the voice of the group is that of an individual person or is generated procedurally. This is that insofar as the group agent is disposed to act for certain fixed or flexible goals across varying scenarios, it must be taken to act on a range of beliefs and other attitudes besides those spelt out by the authorized voice. Let the group endorse the belief that p and that if p, q, in virtue of the dictates of that voice, and it may be expected to be disposed – ultimately, by virtue of an unvoiced belief common to all members – to reason from those premises to the conclusion that q: that is, by at least one account of reasoning, to believe that the premises entail the conclusion (Broome 2013; Pettit 2016). This observation amounts to holding that when we treat groups as agents, as their patterns of behavior require, we inevitably take a step that we also take in the interpretation of individual persons: we fill out the web of attitudes needed to make full intentional sense of what they do.

5. Having a capacity for commitment

Strictly speaking, a group agency might form just on the basis of embodying the four elements listed. But if those elements are in place then inevitably the group agent will be capable of committing itself, via the voice it constructs, to the attitudes it holds. And that capacity for commitment is probably essential for group agents to do the sorts of jobs they are generally designed to do. Those jobs involve cooperation or alignment with

[9]Notice that this means that the members of the group remain ultimately in charge; there is no suggestion that group agency – or group autonomy in the sense explained – deprives members of their power. See (Moen 2023) on related issues.

other agents, as in the relations and exchanges among corporations and between corporations and their creditors, customers and workers. And the capacity for commitment is crucial in building such relationships and in managing such exchanges.

The idea of a committing to an attitude is best introduced by contrast with reporting about an attitude. Suppose I report to you about what I believe or desire or intend, as I might report about another's attitude, or about any state of affairs. If I fail to act according to the attitude I report and you charge me with having misled you I can resort to either of two salient, often plausible excuses; these are explanations that would help to save my reputation as a reliable speaker. I can claim that I got the evidence wrong: that I was misled about my own mind, as I might have been misled about the mind of another. Or I can claim that I changed my mind and had dropped the attitude reported after speaking to you and before acting.

These face-saving explanations are readily available and to that extent my words will be cheap and probably unconvincing. How might I hope, then, to make them more convincing and credible? Presumably by making them more expensive: by staking more on the accuracy of the words I use, exposing myself to a greater cost in the event of not living up to them in action.

I could do this, saliently, either by communicating my attitude in a way that renounced appeal to the misleading-mind excuse or, to take the obvious alternative, in a way that also renounced appeal to the changed-mind excuse. Presumably, linguistic conventions would enable me to indicate that I was doing one or the other. In the first case, to introduce the words as terms of art, I would be *avowing* the attitude rather than just reporting it (Bar-on 2004); in the second, I would be *pledging* the attitude rather than just avowing or reporting it. But how could I have the assurance about an attitude that would enable me to avow or pledge an attitude: to commit to it in either sense?

It is worth nothing that the acquiescence of individuals in going along with others in operating a group agent will be facilitated if each can give others the sort of assurance that pledging their intention to acquiesce would instill.[10] But this is not the place to explore how individual

[10]This observation does not involve taking sides on the divergence between Gilbert and Bratman as to whether joint action requires commitment. The sense of commitment at issue in that debate is essentially normative, whereas here it is not; it is the sense of commitment invoked by game theorists when they say that if you stake money on the accuracy of your communication about an attitude, you commit to the attitude.

human beings may muster confidence enough to make avowals and pledges of their attitudes, although undoubtedly they can (Pettit 2018). The important point to register in the present context is that group agents certainly are in a position to muster that confidence and to avail themselves of the linguistic means of making avowals and pledges.

The voice that speaks for a group, as we saw, is not the voice of a reporter that seeks to track independently existing attitudes. Rather is it a creative voice that, in the context of member acquiescence, makes it the case – and manifestly makes it the case – that the group has any attitudes that it ascribes. But this means that when that voice ascribes a belief or desire or intention, then the group cannot hope to explain a failure to act according to that attitude by the claim that the evidence about the group mind was misleading: the voice wasn't guided by evidence about the group mind, so could not have been misled by such evidence. Thus, the expression of the attitude inevitably has the status of an avowal. It communicates the attitude in a relatively expensive and credible manner, sidelining the possibility of appealing to a misleading-mind explanation of a failure by the group to live up to the words uttered in its name.

But, speaking through a creative voice, the group can go further still. The group not only has the capacity to avow attitudes: this, because its self-ascriptions of attitudes inevitably count as avowals. It also has the capacity to pledge attitudes, or at least to pledge intentions.

Suppose that a group recognizes that there is a large benefit to be won by giving other groups or individual agents – perhaps its own members – the highest possible degree of confidence that it will act in a certain way; in particular, the degree of confidence that would go with a pledge that forecloses the changed-mind as well as the misleading-mind excuse. The promise of this reward can enable it to foreclose the possibility of changing an intention to act in that way, since the foreseen reward – and the reward of proving that its word is its bond – may promise to be enough to keep the intention in place.[11]

In such a situation, the group can use its voice to pledge that it will act in the relevant manner, perhaps conditionally on others acting in a

[11]Such a prospective reward could not enable the group to pledge a belief or a desire: that is, a desire that is premised on the desired state of affairs being attractive under a certain aspect: attractive for the possession of a given desideratum. The group could not foreclose the possibility of changing its beliefs insofar they are designed to reflect data beyond its control; and it could not foreclose the possibility of changing its desires insofar as they are assumed to reflect specific desiderata that lie also beyond its control.

complementary fashion. It can communicate its intention to perform that action in a way that forecloses the possibility of appealing to either a misleading mind or a changed mind in explaining a failure to fulfill the intention. Persuaded of the attraction of eliciting the highest degree of confidence in potential interactants, it can rely on its creative voice, given the individual acquiescence of members in that voice, to make it the case that it will stick with the intention communicated, and to stick with it regardless of factors that might otherwise have led it to a reconsideration of the intention.

These observations are going to be manifest to all individual and group agents, so that organized groups will each be conscious of being able to perform as commissive agents, using active or explicit commitments to establish relations of reliance with one another and with individuals. But group agents will also find themselves having to endorse various virtual commitments in the context of social life. And so, they will be deeply embedded in a commissive, communal web.

Suppose it is a matter of manifest expectation in a society that individual and group agents will generally conform to certain socially salient norms, such acting peacefully, speaking truthfully, keeping their promises and respecting property. In the presence of such matters of common expectation, each will be conscious of others assuming that they are pledged to acting by such norms, and will not try to excuse a failure to conform by reference to misleading-mind or changed-mind excuses. Not saying 'Nay' to that assumption, then, will naturally be taken by all to be equivalent to saying 'Yea'; it will count as endorsing a manifold set of virtual as well as active pledges – and, in relevant cases, avowals.

When I enter a pattern of active and virtual commitments, as each of us invariably does, I effectively hold out an image of myself that I invite others to rely on: speaking for myself, I say 'this is who I am; this is the me you can rely on'. Playing spokesperson for myself, I assume a certain persona, and accept that I must own up to any failures that I cannot excuse – there are likely to be many – and can only apologize to others, in the hope of being forgiven. By an account endorsed by Hobbes in the context of discussing group agency, this is what makes me a person as distinct from just any old agent. I am a person insofar as I can and do personate, as he says, using that term to refer to the self-presentation implicit in the pattern of active and virtual commitment at the core of social life (Pettit 2018).

As this is true for individuals, so it is going to be true for any group agent that individuals constitute, as they rally behind a voice that

speaks for them collectively and acquiesce in its guiding role. Each group agent is going to be a person, then, in the functional sense identified by Hobbes, if not in the intrinsic sense in which personhood is treated as a moral status. Each is going to personate or represent itself via its spokes-voice – the voice that serves in the role of a spokesperson – inviting other individual or group agents to rely on its living up to what it says. And in doing this it can play an effective role in the social world of persons, building up relations of reliance, reciprocation and trust.

Group agents can do this in all sorts of informal ways, as when the CEO promises that a company will look into a complaint, or the head of a church promises that it will explore a proposed initiative. But group agents play a personating role most saliently insofar as they operate within the law, making and managing a range of commitments: the active commitments associated with contract law, for example, or the virtual commitments to abide by ownership rights that are embedded in property law. They undertake commitments, present themselves as aware of the responsibilities imposed by their commitments, and accept that they can hold one another legally responsible for failures to honor them, suing and being sued in court. They operate, as we say, in the manner of legal persons; they function as persons, in the Hobbesian sense of the term.[12]

6. Conclusion

We saw in the first section that no system can count as an agent, and no group can count as a group agent, unless it has a capacity, things going well, to pursue certain goals reliably across different scenarios; to form representations reliably of the scenarios it faces; and to pursue those goals according to those representations. The remaining sections in the paper provide answers to a series of questions that that observation raises.

How can the members, being independent agents, get to act as a group agent's goals and representations dictate? By intentionally, if not always voluntarily, acquiescing in what such action requires of them. How can they identify the goals and representations by which to be guided? By authorizing a voice, like the voice of a single spokesperson, to provide that guidance. But is it possible for them to authorize a

[12]The capacity of group agent to perform as legal persons was already emphasized by medieval thinkers, who focused on the functional profile in law of a body like a professional guild or a monastic order, a town or a parish. See (Canning 1980; Kantorowicz 1997).

voice without relying on a single spokesperson? Yes, by actively constructing the voice to follow: while it may take bottom-up inputs from members – while it may in that sense be responsive to members – it will impose top-down revisions to guard against outputting irrational guidance. Is such an agent capable, finally, of committing like a person to the attitudes the spokes-voice ascribes? Yes, insofar as that voice can avow the group's attitudes, and pledge its intentions, binding itself on pain of sanction to a *persona* on which it invites others to rely.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Bar-on, D. 2004. *Speaking My Mind: Expression and Self-knowledge*. Oxford: Oxford University Press.

Bratman, M. 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.

Broome, J. 2013. *Rationality Through Reasoning*. Oxford: Blackwell.

Canning, J. P. 1980. "The Corporation in the Political Thought of the Italians Jurists of the Thirteenth and Fourteenth Century." *History of Political Thought* 1: 9–32.

Couzin, I. D., and S. A. Levin. 2015. "Special Issue: Collective Behavior." *Journal of Statistical Physics* 158 (3).

Dietrich, F., and C. List. 2007. "Arrow's Theorem in Judgment Aggregation." *Social Choice and Welfare* 29: 19–33. doi:10.1007/s00355-006-0196-x.

Durkheim, É. 1938. *The Rules of Sociological Method*. New York: The Free Press.

French, P. A. 1979. "The Corporation as a Moral Person." *American Philosophical Quarterly* 16: 207–215.

Gilbert, M. 2015. *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press.

Hobbes, T. 1994. *Leviathan*. Edited by E. Curley. Indianapolis: Hackett.

Hölldobler, B., and E. O. Wilson. 2009. *The Superorganism: The Beauty, Elegance and Strangeness of Insect Societies*. New York: Norton.

Kantorowicz, E. H. 1997. *The King's Two Bodies: A Study in Mediaeval Political Theology*. Princeton, NJ: Princeton University Press.

Kornhauser, L. A., and L. G. Sager. 1993. "The One and the Many: Adjudication in Collegial Courts." *California Law Review* 81: 1–59. doi:10.2307/3480783.

List, C., and P. Pettit. 2002. "Aggregating Sets of Judgments: An Impossibility Result." *Economics and Philosophy* 18: 89–110. doi:10.1017/S0266267102001098.

List, C., and P. Pettit. 2004. "Aggregating Sets of Judgments: Two Impossibility Results Compared." *Synthese* 140: 207–235. doi:10.1023/B:SYNT.0000029950.50517.59.

List, C., and P. Pettit. 2011. *Group Agency: The Possibility, Design and Status of Corporate Agents*. Oxford: Oxford University Press.

List, C., and B. Polak. 2010. "Introduction to Judgment Aggregation." *Journal of Economic Theory* 145 (2): 441–466. doi:10.1016/j.jet.2010.02.001.

Ludwig, K. 2016. *From Individual to Plural Agency: Collective Action*. Vol I. Oxford: Oxford University Press.

Moen, L. J. K. 2023. "Eliminating Group Agency." *Economics and Philosophy* 39: 43–66. doi:10.1017/S0266267121000341.

Pettit, P. 1993. *The Common Mind: An Essay on Psychology, Society and Politics*, paperback edition 1996. New York: Oxford University Press.

Pettit, P. 2001a. "Deliberative Democracy and the Discursive Dilemma." *Philosophical Issues* 11: 268–299. doi:10.1111/j.1758-2237.2001.tb00047.x.

Pettit, P. 2001b. *A Theory of Freedom: From the Psychology to the Politics of Agency*. Cambridge: Polity and Oxford University Press.

Pettit, P. 2003. "Groups with Minds of their Own." In *Socializing Metaphysics*, edited by F. Schmitt, 167–193. New York: Rowan and Littlefield.

Pettit, P. 2007. "Responsibility Incorporated." *Ethics* 117: 171–201. doi:10.1086/510695.

Pettit, P. 2008. *Made with Words: Hobbes on Language, Mind and Politics*. Princeton: Princeton University Press.

Pettit, P. 2016. "Broome on Reasoning and Rule-following." *Philosophical Studies* 175: 3373–3384.

Pettit, P. 2018. *The Birth of Ethics: Reconstructing the Role and Nature of Morality*. Oxford: Oxford University Press.

Pettit, P. 2023. *The State*. Princeton, NJ: Princeton University Press.

Pettit, P., and D. Schweikard. 2006. "Joint Actions and Group Agents." *Philosophy of the Social Sciences* 36: 18–39. doi:10.1177/0048393105284169.

Scanlon, T. M. 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.

Searle, J. 2010. *Making the Social World: The Structure of Human Civilization*. Oxford: Oxford University Press.

Seeley, T. D. 2001. "Decision Making in Superorganisms: How Collective Wisdom Arises from the Poorly Informed Masses." In *Bounded Rationality: The Adaptive Toolbox*, edited by G. Gigerenzer, and R. Selten, 249–262. Cambridge, MA: MIT Press.

Tollefsen, D. 2015. *Groups as Agents*. Cambridge: Polity.

Tuomela, R. 2007. *The Philosophy of Sociality: The Shared Point of View*. Oxford: Oxford University Press.