Episteme, 9, 3 (2012) 283–291 © Cambridge University Press
doi:10.1017/epi.2012.17

THE NORMATIVE STANDING OF GROUP AGENTS

RACHAEL BRIGGS
formal.epistemology@gmail.com

ABSTRACT

Christian List and Philip Pettit (henceforth LP) argue that groups of people can be agents – beings that believe, desire and act. Their account combines a non-reductive realist view of group attitudes, on which groups literally have attitudes that cannot be analyzed in terms of the attitudes of their members, with methodological individualism, on which good explanations of group-level phenomena should not posit forces above individual attitudes and behaviors. I then discuss the main normative conclusion that LP draw from the claim that group agents exist: that we ought morally to grant legal rights and responsibilities to group agents, but that group rights should be more limited than individual rights. I argue that when it comes to the fitness of group agents to bear legal rights and responsibilities, LP can draw support from nonreductionist views elsewhere, particularly in the philosophy of mind. I raise some objections to LP's views about the value of granting legal rights and responsibilities to group agents.

In *Group Agency: The Possibility, Design, and Status of Corporate Agents* (2011), Christian List and Philip Pettit (henceforth LP) argue that groups of people can be *agents* – beings that believe, desire and act. The book's key achievement is its development and defense of a new *nonreductive realist* metaphysics for group agents. The view is realist in treating apparent attributions of beliefs and desires to groups as literally true (at least some of the time), and nonreductive in holding that such attributions cannot always be paraphrased into claims about individual attitudes. The view is also *methodologically individualist*: it holds that 'good explanations of social phenomena should not posit any social forces other than those that derive from the agency of individuals: that is, from their psychologically explicable responses to one another and to their natural and social environment' (3). Methodological individualism may require unpacking, but for a contemporary audience, it doesn't require much justification. LP spend their energy arguing for the realist and nonreductionist parts of their view.

LP justify their realism by an extended appeal to functionalism. To have beliefs and desires, they claim, is just to have states that interact appropriately with each other, and with the outside environment. Therefore, if a group is capable of instantiating the right kinds of states, it is capable of believing and desiring. They adopt this functionalist stance not just toward belief and desire, but also toward complex agentive capacities like personhood and self-identification, and functionalism figures heavily in their arguments for group personhood and group responsibility.

LP justify their nonreductionism by claiming that the attitudes of groups supervene on the attitudes of their members, but only holistically – that is, what a group believes about a

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

proposition *p* is fixed by its members' beliefs, but may depend on their beliefs about a wide range of matters besides *p*. Thus, a sentence like 'the US wants to reduce its economic dependence on China' cannot be translated into a claim about the desires of individual US citizens or officials' to reduce US economic dependence on China – instead, it expresses a complicated and disjunctive proposition about a wide range of individual attitudes.

In addition to the development of nonreductive realism, the book makes several other important contributions to the literature. LP draw illuminating parallels between the philosophy of the social sciences and the philosophy of mind. They compare nonreductive realism in the metaphysics of group agents to nonreductive physicalism in the philosophy of mind, discuss an argument against nonreductive realism that parallels the exclusion argument against nonreductive physicalism (more on that soon), and develop a functionalist account of personhood that is structurally similar to functionalism in the philosophy of mind.

LP also provide a clear, user-friendly overview of formal voting paradoxes. These paradoxes present a *prima facie* challenge to LP's nonreductive realism. One way of glossing the paradoxes is this: a group cannot have beliefs and desires that both supervene on the beliefs and desires of group members in a suitably strong fashion and are guaranteed to be coherent under a wide range of counterfactual circumstances. (LP's holism escapes the paradoxes by weakening the supervenience requirement.) Other voting paradoxes raise trouble for the design of group agents: formal considerations make it difficult for group agents to make efficient use of the information available to them, and to avoid manipulation by group members with vested interests. Although LP gloss over some of the technical details – there is only so much room in a single book – they sketch the general structure of the formal results, explain their significance and summarise possible avenues of response. The literature on all these topics is tangled and technical, but LP's overview is admirably accessible.

The book has one important limitation: it is an argument that group agents are possible, not an argument that group agents exist. While LP draw on facts about history and law, their methods are chiefly *a priori* and analytic. Although it seems overwhelmingly likely that group agents exist to fill the conceptual space that LP have carved out, we will need more investigation into the actual dynamics of group reasoning to establish which agential capacities groups have, and how group attitudes relate to the attitudes of group members.

In the remainder of this review, I will focus on the normative consequences of LP's view. Suppose there are group agents, as LP claim. What do these agents owe to the rest of us, and what do we owe them? In Part III of the book, the authors address these questions. Their discussion involves two kinds of normativity, which we should explicitly distinguish. LP take themselves to be discussing the *moral* justification of *legal* and other societal norms. When they discuss rights and obligations, they always mean conventional rights and obligations – the type of rights and obligations we invoke when we say I owe money to the telephone company, or I have a right to Worker's Compensation. LP are interested in the question: which systems of legal norms are morally best or right? In particular, is it morally appropriate for a system of social conventions to grant rights and responsibilities to group agents?

LP's way of formulating their question quarantines it from the debate between consequentialists and deontologists in normative ethics – everything they say is compatible both

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

with consequentialism and with deonotology. Likewise, LP's formulation of the question is neutral with respect to exactly how the true moral theory grounds the right/best social arrangement (or arrangements, if there is more than one). The many consequentialist options include: the best social arrangement is the one that maximises expected value; any social arrangement is right so long as it satisfices expected value; the best social arrangement is the one that maximises not expected value, but value full stop. The many deontological options include: the right social arrangement is the one in which conventional rights correspond most closely to natural rights; the right social arrangement is the one that everyone would agree to in a state of nature; acceptable social arrangements are ones in which the government does not infringe on any individual's natural rights.

Whatever the relationship between the true moral theory and the best social arrangement, however, our conventions for dealing with group agents must somehow be justified on moral grounds. And the relevant moral considerations must be based in what is good or bad for agents – whether we understand goodness and badness in terms of pleasure and pain, autonomy and rights-violations, desire satisfaction and desire frustration, or in some other way. But whose interests matter?

LP adopt an assumption of *normative individualism*, which they gloss as 'the view that something is good only if it is good for individual human, or more generally sentient, beings' (82). The idea is that that only individuals' interests are morally significant, so that granting political rights and obligations to groups is justified only insofar as it benefits individuals.

LP say almost nothing in support of normative individualism. There are good dialectical reasons for this: normative individualism is extremely uncontroversial. Nonetheless, I believe LP ought to say more. The take-home message of the book is that groups can resemble individual agents much more closely than we might think: like individuals, groups can believe and desire contents, carry out actions, be held responsible for what they do, bear rights, perform as persons, and entertain first-person indexical thoughts about themselves. So it is natural to ask whether the resemblances extend to the moral realm.

In order to assess normative individualism, we must determine whether anything can be good or bad for groups. If groups are capable of benefitting or suffering in ways that would be morally relevant in individuals – if groups can feel pain, or enjoy autonomy – then it seems that we are morally obligated to attend to the interests of groups.

To address this question, we'll need a quick look at the concept of wellbeing – what constitutes a benefit or a harm. Following Parfit (1984: 493), we can divide accounts of wellbeing into three main camps.

- *Hedonic Theories* hold that experiences are good for an agent insofar as they are pleasurable, and bad for the agent insofar as they are painful.
- *Desire-Fulfillment Theories* hold that things are good for an agent insofar as they fulfill the agent's desires, and bad for an agent insofar as they frustrate the agent's desires.
- *Objective List Theories* hold that certain things (such as knowledge and friendship) are good for all agents, regardless of whether the agents desire them.

Which of these approaches, if any, make it likely that groups have morally commanding interests?

First, consider hedonic theories. Although LP have argued that groups have beliefs and desires, it seems unlikely that groups experience pleasure or pain. But even here, there is

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

some cause for concern. On a sufficiently expansive conception of functionalism, pleasure and pain are functional kinds: anything with states that play the pain role experiences pain. And while ordinary groups don't seem to have states that play the pain role, it's not so obvious that there *couldn't* be a group that felt pain. So if hedonic theories of well-being are true, then normative individualism is probably correct, but may be only contingently correct.

On desire-fulfillment theories, it is impossible for anyone with LP's commitments to avoid the conclusion that groups have interests. Given that there are group agents, there are groups with desires, and it follows immediately that groups have interests.

Objective list theories stand somewhere between hedonic theories and desire-fulfillment theories. Some putative goods, like physical health, are available only to individuals – it makes no sense to speak of the physical health of a group, over and above the physical health of its members. Another putative good, knowledge, is presumably available to any agent capable of the right functional states. Still others occupy a puzzling middle ground: can groups act autonomously, have friends, or perform meaningful work? However, it looks as though at least some objective list theories will count groups as having morally commanding interests.

So on a variety of views about wellbeing, LP's realism about group agents calls normative individualism into question. Let us move from this theoretical issue to LP's practical prescriptions. LP argue that it is morally best or right to endow groups with both conventional obligations and conventional rights. However, they hold that there is an important asymmetry between obligations and rights: while groups should be held responsible for their actions to the same extent as individuals, they should not have the same rights as individuals.

LP's argument in favor of group rights and obligations can be divided into two basic components. The first establishes that that rights and obligations can meaningfully apply to groups – that treating groups as parties to the social contract need not be misguided, a useful fiction or an honorific practice. The second establishes that treating groups as bearers of rights and responsibilities is beneficial – which, given LP's normative individualism, means beneficial to individuals.

First, let us consider LP's argument that groups are fit to bear rights and responsibilities. LP split the argument into two parts, arguing in chapter 7 that groups can be held responsible for their actions, and in chapter 8 that groups can be persons (and hence bearers of rights).

In chapter 7, LP claim that the following three conditions are individually necessary and jointly sufficient to make an agent fit to be held responsible (155).

- *Normative Significance*: The agent faces a normatively significant choice, involving the possibility of doing something good or bad, right or wrong.
- *Judgmental Capacity*: The agent has the understanding and access to evidence required for making normative judgments about the options.
- *Relevant Control*: The agent has the control required for choosing between the options.

Since the three conditions are both uncontroversial and plausible, LP devote the bulk of their energy to arguing that groups are capable of meeting the conditions.

To meet the normative significance requirement, a group need only be an agent and to encounter the right kinds of choices. If there are group agents, their everyday interactions with others will doubtless furnish them with normatively significant choices.

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

Judgmental capacity requires more than agency: it requires that the agent be capable of reflecting on its own beliefs and desires. As LP point out in chapter 2, a normative belief that *q*-ing is good is not merely a desire to *q*. Very simple agents – robots and animals – can have desires without having any normative concepts in which to couch normative beliefs. Still, desire and normative belief occupy similar functional roles: both someone who wants to *q* and someone who believes that *q*-ing is good will *q* when given the chance (all other things being equal, and provided they have suitable beliefs). So what's the difference between desiring something and believing that it's good?

LP's answer in chapter 2 provides little concrete guidance. They suggest that an ascription of a mental state to an agent should neither give rise to expectations that the agent can't meet ('the ceiling constraint'), nor fail to give rise to expectations that the agent can meet ('the floor constraint') (29). Both principles sound sensible enough, but which expectations can we use to distinguish desires from normative beliefs? (The same problem arises for other pairs of sophisticated and simple states: what exactly is the difference between a belief that *p* is probable and a high credence in *p*, or a belief that *p* entails *q* and a disposition to infer *q* from *p*?)

Luckily, in chapter 7, LP offer more help. They write, 'A group forms a judgment or other attitude over a certain proposition when the proposition is presented for consideration . . . and the group takes whatever steps are in its organizational structure for endorsing it' (159). In other words, we can make fine-grained distinctions among the contents of the group's beliefs by looking at which sentences the group uses to express its commitments. (This suggests an answer to the more general puzzle in chapter 2: we can use linguistic behavior to break ties between attitudes with otherwise similar functional roles. For an agent to count as having a sophisticated attitude – such as a belief that *p* is good, a belief that *p* is probable, or a belief that *p* entails *q* – it must be capable of producing the right sorts of sentences at the right times.)

So a group has the appropriate judgmental capacity when its members can use its organisational structures to form collective judgments about the normative significance of its actions. Notice that the group needn't actually form any normative beliefs in order to satisfy the judgmental capacity condition. It need only be capable of forming the appropriate normative beliefs. LP suggest that there is no good reason to form groups whose organisational structure prevents deliberation about the normative significance of its decisions – in fact, to form such a group would be negligent. Therefore, in a sufficiently just society, group agents that satisfy the normative significance requirement will also satisfy the judgmental capacity requirement.

What about the requirement of relevant control? Here, LP face a significant obstacle: there is a plausible argument for the conclusion that groups never have control of their actions. I quote LP's formulation of the argument on p. 160.

1. Whatever a group does is done by individual agents.
2. Individuals are in control of anything they do, and so in control of anything they do in acting for a group.
3. One and the same action cannot be subject both to the control of the group agent and of one or more individuals.

Therefore,

4. The group agent cannot be in control of what it does; such control rests exclusively with the individuals who act for the group.

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

LP respond to the argument by denying premise 3. They point out that, in many domains, a single event may have causes on more than one level at once. The most famous – but also the most controversial – of these domains is psychology, where the same action can arguably be caused by both by a mental state and by the physical state that realises it. LP use the example of a flask in which water is boiled, and which consequently breaks. Intuitively, the breaking of the flask is caused both by the boiling of the water and by the individual water molecule that hits the side of the flask with the right position and momentum to trigger a break. The two causes are of different types: the boiling of the water is a 'programming cause', while the impact of the individual molecule is a 'triggering cause'.

In distinguishing between programming causes and triggering causes, LP appeal to an account of higher-level causation defended by Jackson and Pettit (2004). But they mention in passing an alternative account due to List and Menzies (2009). This alternative account deserves attention, since it provides interesting new grist for LP's mill. In particular, it suggests that we should reject premise 1 of the argument as well as premise 3.

List and Menzies hold that causation is difference-making, cashed out as follows.

> The presence of F makes a difference to the presence of G in the actual situation if and only if (i) if any relevantly similar possible situation instantiates F, it instantiates G; and (ii) if any relevantly similar possible situation instantiates ~F, it instantiates ~G.[1]

They spell out conditions (i) and (ii) formally in counterfactual terms, then use the logic of counterfactual conditionals to prove that, under appropriate circumstances, the presence of G can be caused by both the presence of a higher-level property F1 and the presence of a lower-level property F2 that realises F1. For LP, the moral is that, under appropriate circumstances, the same action can be controlled both by the group and by its individual members.

But List and Menzies's difference-making conception of causation has another consequence: it allows higher-level causes to trump lower-level causes. So it may be that certain mental states cause actions while their neural correlates do not, that boiling water causes the breaking of the flask while the triggering molecule does not, and that groups control actions that are not controlled by any individual. In general, a higher-level property F1 trumps a lower-level property F2 as a cause of G iff F1 is a difference-making cause of G and G is present in some closest ~F2 worlds that are F1 worlds. In the special case at hand, this means that a state of the group, and not a state of the individual actor, controls an action of the group iff the action depends counterfactually on the group's state, but might have occurred even if the individual state giving rise to the action hadn't.

The upshot is that, under some circumstances, individual actors are not in control of the group's actions. In other words, we must deny either premise 1 or premise 2. I suggest that premise 1 is to blame: though individuals always control their own actions, some group actions are not identical to any individual action. Rejecting premise 1 might

1 There is a slight disconnect between List and Menzies's framework and LP's framework. List and Menzies discuss relations of causation between instances of properties, while LP discuss relations of control between agents and action. I'll assume that sentences of the form 'agent S controlled action A' can be translated into sentences of the form 'agent S's instantiating mental state F caused the presence of property G', where G is present iff A takes place.

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

seem dubious: won't it require a commitment to mysterious social forces that operate independently of individual actions? No: I claim that rejecting premise 1 is perfectly compatible with methodological individualism.

On the picture I propose, an individual's action on behalf of a group bears the same relation to the group's action that the exact microphysical state of the water bears to the water's boiling, or the physical state of an individual's brain bears to the physical state of their mind. The action of a group agent – for instance, the mailing of promotional materials by a lobbying group – can be realised by more than one individual action – the material might be mailed by any of three secretaries in the group. The actions of groups supervene on the actions of their members, and are realised by the actions of their members, but are not identical to them.

My amendment provides further support for LP's claim that groups can meet the relevant control condition. Thus, LP are reasonable to assert that group agents satisfy – or could be designed to satisfy – all three conditions for fitness to be held responsible. It need not be a mistake, a metaphor or a fiction to hold groups responsible for their actions. Does it also make sense to attribute rights to groups?

LP frame their discussion of rights in terms of the concept of a *person*. They never give an explicit, theory-neutral definition of the term 'person', but we can triangulate the meaning of the term based on the purpose it serves in the argument: a person is the sort of thing to which it is appropriate to assign conventional rights. A society can endow human beings can with property rights and voting rights, but it cannot endow rocks or trees with any rights at all. Human beings are persons; rocks and trees are not.

LP favor a *performative* conception of personhood, on which

> To be a person is to be able to perform as a person. And to perform as a person is to be party to a system of accepted conventions, such as a system of law, under which one contracts obligations to others and . . . derives entitlements from the reciprocal obligations of others. In particular, it is to be a knowledgeable and competent party to such a system of obligations. (173)

In other words, to be a person is to be capable of the understanding and competence required for participation in a social contract. This understanding and competence make it appropriate to grant conventional rights to an agent. Or in slogan form: the social contract is a game open to anyone and anything capable of playing.

LP contrast their performative conception of personhood with a rival *intrinsicist* conception, on which 'there is something about the "stuff" persons are made of that distinguishes them from non-persons: something that makes persons stand out' (171). Phrased this way, the intrinsicist conception sounds mysterious, and LP say little to remove suspicion. Having the right intrinsic properties is presumably a matter of being made from the right type of ectoplasm or protoplasm – a dubious picture at best.

Given the weakness of the intrinsicist conception, it's no wonder LP embrace its performative rival. But there is another more attractive conception of personhood that LP don't consider, which I'll call *interest-based*.

According to the interest-based conception of personhood, a person is the kind of creature for which things can go well or badly. The appeal of the interest-based conception is this: it seems pointless to grant rights to something that cannot benefit from those rights. That is why the idea of animal rights makes sense: even if animals cannot navigate the human social contract, they can still suffer harms, and so it makes literal sense to develop

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

social conventions on their behalf. Likewise, that is why the idea of computer rights make less sense, insofar as it is unlikely even a very clever computer could benefit from social protections.

LP may have methodological reasons for ignoring the interest-based conception of personhood. Throughout the book, they take care to distinguish between descriptive questions – what makes a group capable of accruing conventional rights and responsibilities? – and normative questions – is it morally good that groups should have conventional rights? But the interest-based conception blurs the distinction; the idea of interests is already a normatively, perhaps morally, loaded concept. Nonetheless, the interest-based conception is appealing enough to deserve consideration.

Given LP's performative conception, however, it seems clear that group agents are capable of holding rights as well as responsibilities. LP argue in chapter 9 that groups can reflectively self-identify – that is, a group can think of itself as one agent among many, as participation in the social contract requires. The group *identifies* with its desires – LP's way of denoting the minimal kind of self-conscious necessary for agency – whenever its members deliberately act on behalf of the group. It *self-identifies* – LP's way of denoting a more sophisticated kind of self-consciousness necessary for participation in a social contract – whenever it identifies in a particular, linguistically mediated mode. (Notice that the relationship between self-identification and identification is roughly the same as the relationship between believing p is good and desiring p, or the relationship between believing p is probable and placing a high credence in p.)

LP hold that first-person pronouns are crucial to self-identification: an individual self-identifies when she uses the pronoun 'I' both in a special first-personal way which does not permit her to be uncertain about whether she is its referent, and in a third-person way that picks her out as one agent among many. And a group self-identifies when it uses 'we' in both of these ways. While LP's exact way of framing their point is subject to linguistic quibbling – perhaps the speakers of a language that allowed pronoun-dropping could form a self-identifying group even if they never employed any pronouns whatsoever – the basic point is sensible. In order to self-identify, a group must think of itself in the right self-conscious, first-personal way, whether or not this type of thought essentially involves pronouns.

So LP make a strong case for the appropriateness of assigning conventional rights and responsibilities to group agents. When we say that the federal government has a right to a certain percentage of my income, or that the electric company is responsible for last night's power failure, we are not making a category error. But is it useful, as well as apt, to assign rights and responsibilities to group agents? LP hold that it is.

Holding groups responsible for their actions, they claim, is a good way of addressing 'deficits of responsibility'. The holistic dependency of group actions on individual beliefs and desires means that groups can intentionally do things none of their members want or approve of. In such cases, a group agent may commit a wrong for which none of its members are responsible. Holding the agent responsible prevents the wrong from going unpunished. And holding a group agent responsible won't prevent us from holding its members responsible, when appropriate. Multiple agents can be responsible for the same wrong, so we should not expect any 'conflicts of responsibility' to arise.

With rights, things are different. LP argue that, in an ideal society, group persons will have rights, because endowing them with rights benefits individuals. Allowing group agents into the social contract lets them participate in a 'regime of mutual respect' in

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press

which all agents interact in a manner that is free from violence, coercion and manipulation. Ultimately, the beneficiaries of this regime are individuals, who gain autonomy because they are respected by group agents.

But LP hold that group agents should not have the same rights as individuals. Groups are vastly more powerful than individuals, and endowing them with equal rights would be a recipe for exploitation. Unlike responsibilities, rights can and frequently do conflict.

In defending these practical conclusions, LP have presupposed normative individualism. If normative individualism is false, how seriously does it damage LP's conclusions? Not too seriously, I suggest. The argument for endowing groups with responsibilities remains the same as before. If groups have morally commanding interests, then we have even more reasons to include them in the social contract: then they too can benefit from a regime of mutual respect.

Still, there are reasons to limit the rights of groups. The fact remains that groups are vastly more powerful than individuals, and there are legitimate reasons to give less powerful individuals special entitlements. The rich in America do not have the same rights to welfare or free school lunches as the poor, and this is as it should be. Some benefits can be provided exclusively to the poor without any morally serious harm to the rich, who already have ample means of getting what they need. Furthermore, even if groups have morally commanding interests, they are unlikely to have the same needs as individuals. It makes sense to outlaw the torture of human beings, but there is no such thing as the torture of corporations.

In summary, then, LP have developed a compelling metaphysics and political philosophy of group agents. They have given us reasons to believe not just in group agents, but in group persons: beings that are fit to be held responsible and fit to be endowed with (limited) conventional rights.

REFERENCES

List, C., and Menzies, P. 2009. 'Nonreductive Physicalism and the Limits of the Exclusion Principle.' *Journal of Philosophy*, 106(9): 475–502.

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

Jackson, F., and Pettit, P. 2004. *Mind, Morality, and Explanation: Selected Collaborations*. Oxford: Oxford University Press.

Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.

RACHAEL BRIGGS is a Research Fellow at the Australian National University and Griffith University. Her research interests include epistemology (particularly formal epistemology and reliabilist accounts of knowledge), metaphysics (particularly truthmaking and the metaphysics of chance) and decision theory.

https://doi.org/10.1017/epi.2012.17 Published online by Cambridge University Press