https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

WHAT IS IT LIKE TO BE A GROUP?

By David Sosa

> This group-think of "we all sat around and decided"—there's only one person that can decide, and that's the president.
>
> —George W. Bush[1]

I. Introduction

A stark distinction between consequentialist and Kantian ethics is that, according to the latter and not the former, the actual consequences of an action are of no significance for ethical status. So these two sorts of ethical theory differ over the relevance of *downstream* phenomena: consequences of the action. I want to investigate how these theories might differ too over *upstream* phenomena: what the action is a consequence of. I will focus on the case of group action and collective responsibility. To preview, my argument will be that there is a kind of analogue to the problem of aggregating the value of consequences—about aggregating responsibility with respect to the roles of cooperating agents—that Kantian theories will not confront and consequentialist theories will. I do not think this cuts decisively in favor of either kind of view; rather, the issue seems to me a useful way to characterize a deep difference between Kantian and consequentialist theories. In fact, consideration of the issue and of the difference it clarifies will point, ultimately, toward a way of making those views compatible: deontological theories are *person-centered,* while consequentialist theories are *act-centered.*[2]

Let me indicate, here, that reflective equilibrium is the method of this essay. I do not have an argument for my position that proceeds deductively from indubitable principles to the conclusions accepted. Instead, I am putting together, into a structure that exhibits some order and exposes some important relations, what seem to me to be attractive views. Even if one does not accept all of the views that I seek to combine, the resulting structure may be sufficiently compelling to attract those who already accept some of the views involved. To those who begin by rejecting all of the views I attempt to combine, the effort will likely appear more or less to beg the question.

[1] Quoted in the *New York Times,* September 2, 2007.

[2] I will not trace the implications of the issue for the position of virtue ethics or of other views relative to consequentialist and deontological theories.

doi:10.1017/S0265052509090098

© 2008 Social Philosophy & Policy Foundation. Printed in the USA.

II. Normative Status Groups Cannot Have

An important issue in the literature on collective responsibility can be understood as one about the *reducibility* of the collective's responsibility to that of the individuals constituting that collective—about the reducibility of the collective as a specifically ethical individual. I will not deeply enter that debate in its own terms. But I will accept some of the presuppositions of at least some of the literature: in addition to being susceptible to harm, sometimes groups *do* good or bad things, and are, accordingly, responsible for various ethical benefits or costs. Like ordinary individuals, then, groups can be both the "patients" of morally significant effects and the agents of those effects (in a sense that leaves open the issue about the reducibility of responsibility). Indeed, unlike George W. Bush, I would allow that groups *can* and *do,* and indeed often *should,* be the ones making the decisions, rather than any one person, even if he is president. But I also think the quotation that stands at the opening of this essay is compelling (in one way or another) because there is a kernel of wisdom in its vicinity—in effect, that is what I will argue here.

Now, a traditional problem for consequentialist theories is how to aggregate harmful effects. To take just one instance of the general form: If each of an agent's alternative actions has indefinitely many harmful effects, then how can any one of those alternatives be better or worse than any other? I have resisted the idea that this phenomenon is a deep or serious problem for consequentialism;[3] but it does not, in any case, arise in the same form in the context of a deontological theory. Let us consider a kind of obverse question: If a number of people perform an action jointly—lifting a heavy piano, say, or encircling a city to which they are laying siege—then how to mete out responsibility? No *one* of the individuals involved is, by herself, responsible for the effect. The group as such is responsible—the individuals are responsible *as* part of the group. I think this problem is, in fact, of the same form as the traditional problem for consequentialist theories, the problem about aggregating harmful effects.

The general form of the problem, I believe, has to do with finding ethically significant order in a relevant complex. In the case of the traditional problem, even when each of a pair of alternatives has indefinitely many harmful consequences, so that in some sense the *sum* of the harm produced by either is equal, sometimes one of the alternatives is ethically better than the other. The challenge is to find an ethically significant order in the complex constituted by the effects and to ground the superiority of the better alternative in its terms: evidently, the order provided by *summation* is inadequate. Of course, this challenge can be met in various ways.

[3] See David Sosa, "Consequences of Consequentialism," *Mind* 102, no. 405 (1993): 101–22.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

The point here will be more attractive if you view, as I do, what might be called[4] the "nearest and dearest" objection to consequentialist theories (which challenges them to account for any preferential treatment of individuals to whom an agent might stand in special personal relations) as another closely related issue, akin to what might be called the "should the numbers count?" problem discussed above for consequentialism.[5] Again, the general issue is that it can seem that consequentialism leaves no room for what seem intuitively to be ethically relevant *articulations* of the space of consequences. We are not concerned merely with what has the best consequences (in any simplistic interpretation thereof), we are concerned also with *how* those consequences are distributed, for perhaps various *how's*. That every act has indefinitely many consequences should not flatten out residual ethical differences, and whether the consequences are distributed in a way that favors one's own family, for example, seems to be of moral moment, one way or another.

To understand how the question about joint action might be viewed as another instance of that sort of problem (about an ethical articulation of a space of causal relations), begin by noting that the "is an effect of" relation has an image in the "is a cause of" relation. Just as we think the consequences of an act are of moment for its ethical status, we think the causes of an act determine the issue of responsibility. The effects of the various actions performed by the ordinary individuals constituting the group are themselves only some of the forces at work leading ultimately to the relevant effect. In fact, there might be indefinitely many forces leading to whatever effect any group ultimately produces. But if in any case of group action there are indefinitely many causes producing the (say) harmful effect, then how to distinguish any one of them as differentially ethically significant? How to select even from among the subjects that are part of the causal history those that are ethically relevant? And then how to select the *group's* joint agency, in particular, as a distinctively ethically responsible subset of the complex total cause? The concept of *causation* cannot, by itself, make that selection.

Again, I do not think this sort of problem cuts deeply against consequentialist theories. It provides insightful challenges for development of the theory: demands for more specific and more adequate accounts of how the consequences produced by one or another act are to be compared, so that the ethical superiority of one of the acts over the other is established according to the theory. Even after crucially important first steps—like noting that the relation between the effects is not *additive*[6]—there is more work to be done. But there are models available to allow

[4] Indeed, it is so called by Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101, no. 3 (1991): 461–82.

[5] The expression is from John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6, no. 4 (1977): 293–316.

[6] Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989).

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

further development. Mundane observations—for example, that although all line segments have indefinitely many points, some are nevertheless *longer* than others, and that sometimes the winner of a competition wins fewer points and loses more than her opponent[7]—lead in useful directions for responding to these challenges.

All of this is controversial, granted. But my purpose here is not specifically to make a convincing case that a particular view of these matters is correct. It is, rather, to highlight how Kantian theories sidestep the issue about meting out responsibility to the individuals involved in group actions in something like the way that those theories circumvent issues affecting consequentialist theories about finding ethically significant order in the consequences.

Why do Kantian theories not confront this issue about aggregating responsibility across a collective? I think it is because groups do not have the sort of *will* required—the sort of will the goodness of which Kant takes to be definitive of ethical status. Groups *can* rightly be said to want things (at least in a theoretically neutral sense), to have intentions, to have beliefs and other states characteristic of a mind. Thus, I cannot deny that there is *a* sense in which a group can have a mind or a will. But I think there is another sense in which a group cannot have a mind or a will and thus cannot be good in the way Kant characterized.

To begin with, a group has a mind only in a derivative sense: the persons that are members of the group have minds, and the group's mind (in whatever sense it has one, its beliefs and desires) is some sort of construct from those minds. Certainly, there seems to be an asymmetric dependence relation of the group's mental states on those of its members. The group has, for example, a desire constitutively in virtue of the states of mind—not necessarily themselves desires—of some of its members. Whatever we might think of the ultimate fundamentality of an individual person's having a mind, it is not plausible that it is always constitutively in virtue of a group's having various states of mind that any of its members has a desire.

This view is supported further by considerations involving punishment. Whether or not it is ultimately defensible to punish a person by making him suffer pain, there doesn't seem to be any special *additional* normative problem about punishing a person by making a proper part of that person suffer harm. Jones is to be punished: the punishment is to be applied by hurting a specific part of his body. There is no temptation, in this case, to find this unfair or inappropriate. It would be a bad joke to say, "But you didn't hurt the rest of his body, and you didn't hurt his feelings!" But the analogous circumstance with groups seems very different. Here, if the group is punished by making a par-

[7] This can be true in tennis, for example, if one wins a set 6-4, having won six games after deuce and lost four at love.

ticular proper part of that group suffer pain, the situation seems unfair and inappropriate.

A thief steals using his right hand: there is nothing problematic (aside from whatever general problems about corporal punishment there may be) about making his left hand rather than his right hand suffer. A corporation steals using its accounting department: if its manufacturing department is sent off to jail (while the accounting department is left alone), one feels that an injustice has been committed. Punishing an individual by punishing (only) a proper part seems unexceptionable, whereas punishing a group by punishing (only) a proper part seems wrong. This supports the idea that groups have certain sorts of normative status only in a derivative way. Again, this is in no way to deny that groups *have* normative status or even to deny that groups have other kinds of normative status in a nonderivative way. But there is an important difference: there is no special problem about punishing an individual person by harming just a proper part, but punishing a group by harming just a proper part of that group seems unacceptable.

The situation is, of course, complex: in the case of a group, the relevant parts are themselves normative subjects; not so in the case of a single individual. In the group case, then, two things are being punished, the group *and* the particular members of the manufacturing department, while in the case of the single individual, the harm that befalls whatever specific part of his body does not constitute a *punishment* of that part. One might seek to explain the difference in our intuitions about punishment by appeal to this difference, and I would not resist. But that is effectively to grant my point about the asymmetric priority of the normative status of individuals over groups. Crucially, it seems inappropriate to harm a group by harming an individual who should not be punished; in contrast, if by punishing an individual who should be punished, one thereby punishes a group that does not (as a group) deserve punishment, we do not have the same feeling that an injustice has been committed.

This distinction in fundamentality (again, groups have minds, if at all, *in virtue of* their constituents' having them—but not vice versa) is, I believe, coupled with a distinction in varieties of normative status to which persons and groups are susceptible. I think it may be useful to see this distinction in terms of a nonethical kind of normativity first: Take a person who believes that p and also believes that q. Whether or not there is any demand on such a person to have a conjunctive belief that (p and q), there does seem to be a demand that she not have the negative conjunctive belief that (not: p and q). But it is hard to see the group as subject to the same demand in the same way: whatever is wrong with a group that believes (whatever it takes for the group, as such, to have beliefs) all of p, q, and (not: p and q), it is not quite the same thing as is wrong with a single person who combines those beliefs, or so I will argue. An attractive explanation of this difference is, first, that the sense in which a group has

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

a mind does not entail that there is a unified locus of consciousness, and, second, that there is a special set of normative features applicable to unified centers of consciousness.

Another useful example is the phenomenon of multiple personality disorder, also known, interestingly, as "*dissociative* identity disorder." Ignoring for a moment the controversy in the medical profession about the status of this diagnosis, the disorder seems to involve what we conveniently treat as a single *person* having more than one *personality.* There is more than one personality, and there is no single unified locus of consciousness. This seems to go naturally with some indeterminacy in how to view such people normatively. One is also tempted to view victims of the disorder as effectively more than one person. I think groups are in many ways similar to people with dissociative identity disorder (again, assuming there is such a thing). As with such individuals, groups are more or less associative and seem to be more subject to every variety of normative evaluation the more substantial the relevant associations among their elements.[8]

These heuristic analogies provide a model for a view on which groups, because they do not have a certain sort of mind (a mind with a single unified locus of consciousness), are not apt for a certain variety of normative evaluation. So I am urging that three phenomena go together: (i) groups are *derivative* bearers of normative status; (ii) groups cannot have a single unified locus of consciousness (and thus cannot have a *will* in one important sense); and (iii) groups are not subject to deontological normative evaluation.

Why do I think these three phenomena go together? Why not think, instead, that although groups are derivative bearers of normative status, and notwithstanding the fact that they lack a single unified locus of consciousness, nevertheless, groups are subject (derivatively) to (even) deontological normative evaluation? Sometimes a group acts for a reason, and sometimes the reason for which a group acts is that the act is seen by the group as its duty: it conceives the act as its moral obligation. Accordingly, sometimes, it seems, a group can satisfy the categorical imperative. Other groups—criminal conspiracies, say—act for different sorts of reasons, because of various *inclinations,* inclinations that could not be seen as binding on any other group undertaking an action. Again, these group inclinations may be best understood as deriving from various other states of mind (perhaps, but not necessarily, themselves inclinations) of the members of the group. But can't the group nevertheless be subject to deontological evaluation, even if only derivatively? This is a serious challenge.

Of course, groups do act for reasons, some good, some bad. Groups can satisfy the categorical imperative or violate it. So if I am right that there

[8] Mark Greenberg usefully suggested the examples of zombies and sleepwalkers: it is plausible that complications about their moral status are related to their not being conscious.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

is a kind of deontological normative status that groups cannot exhibit, then the view I defend may amount to a development of the Kantian line. On the view defended here, satisfaction of the categorical imperative is not sufficient for the sort of normative status Kant sought to characterize: that status is limited to entities that are themselves conscious. The fact that something is constituted by parts that are conscious does not mean that the complex is itself conscious (though we might say that the thing "involves consciousness"). The problem is not that the complex doesn't have *enough* consciousness: two instances of consciousness is as bad as none at all—it is "one center of consciousness too many," to torture Bernard Williams's nice phrase.[9] When counting subjects, we insist on a one-to-one correspondence with unified centers of consciousness. Compare again individuals with multiple personality disorder: shall we now say that in addition to the various centers of consciousness "inhabiting" the one body, there is also yet another center of consciousness involved, the one built up from all of those others?

When we think about the natural attitude toward such cases (and cases of severe schizophrenia may be others in point), I think the line I am defending clarifies. *Sybil,* the 1976 made-for-TV movie (based on a book by Flora Schreiber), dramatized the life of Shirley Ardell Mason, a psychiatric patient who underwent Freudian psychotherapy in the 1950s with Dr. Cornelia Wilbur. According to Wilbur, there were sixteen personalities in some sense "within" Mason. The movie's dramatic interest, I believe, is largely the result of our complicated attitudes toward Ms. Mason's status. We are almost unavoidably drawn to treating her as a single person.[10] But we are sensitive to the fact that the disintegration of her consciousness demands a more subtle attitude.

The analogue here of my view about groups would be that Ms. Mason (on the assumption that the claims about her were accurate), while certainly a complex individual capable of action and apt for one sort of normative evaluation, is not apt for another sort of normative evaluation. Sure, in an important sense, Mason could have been worse off. She could have lacked the leavening effect of her more reasonable parts. But the sort of evaluation we make of Mason is not at all of the same sort as that we make of ourselves when we have trifled with a person's feelings, or been delighted by their ill-fortune, or been pursuing things for our own illegitimate gain, at great cost to others. While there is something it is like to be any one of Ms. Mason's various personalities (some were shy and passive, others bolder and more aggressive, and so on), there is nothing it is like to be Ms. Mason "herself." So even if Ms. Mason acts intentionally, perhaps because one of her

[9] Bernard Williams, "Persons, Character, and Morality," in Amélie Rorty, ed., *The Identities of Persons* (Berkeley: University of California Press, 1976), 214.

[10] Cf. Bernard Williams, "Personal Identity and Individuation," *Proceedings of the Aristotelian Society* 57 (1956–57): 229–53.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

personalities is able to persuade or dominate the others, her action cannot be seen as the expression, in the relevant sense, of a good will. Suppose that Mason's reason (*not* that of the relevant one of the various persons in her) for acting as "she" does is that a given act is her duty. Even so, we should, I think, resist the claim (in at least one of its interpretations) that *Mason* ("herself," as we are inaptly tempted to say) is good: there is no single, sufficiently unified person there to be good. If this is the reasonable thing to say about such cases, then I think we should say the corresponding thing about groups. An interesting result—perhaps a bit surprising and not welcome to all, but not, I think, immediately out of the question—is that *consciousness* is a presupposition of the sort of normative status distinctive of deontology, one that contrasts with consequentialist evaluation.

I think a distinction from the philosophy of mind comports well with the picture developing here. Familiar functionalist views about propositional attitudes can be seen as intuitively less plausible in application to phenomenal states.[11] One possible view in the philosophy of mind would hold, accordingly, that although various nonbiological objects could satisfy the functional organization characteristic of having cognition, our intuitive resistance to the result that such objects would be thinking things is grounded in a presupposition that only conscious entities can be thinking things—where being conscious is coeval with having *phenomenal* states, and where that is *not* understood in functionalist terms (or perhaps understood at all).

In any case, for what it's worth (I am not in fact myself attracted to functionalist conceptions of propositional attitudes), the view developed here would fit well with that sort of line. A thermostat in some sense "notices" that the temperature has risen to 75 degrees Fahrenheit; that's why it "does" its thing (switching on the compressor). The naturalness of this way of speaking is understood in terms of the (modest) similarity between the relevant functional relations instantiated in the thermostat and those characteristic of human belief and action. But our resistance to a literal reading of that natural way of speaking about the thermostat is due to the fact that the thermostat is not conscious: there is nothing it is like to be a thermostat. And there being something "it is like" is perhaps made most salient in terms of phenomenal states.

David Hume said, "For my part, when I enter most intimately into what I call myself, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch myself at any time without a perception, and never can

[11] See, e.g., Jaegwon Kim, *Mind in a Physical World* (Cambridge, MA: MIT Press, 2000). I myself find the inegalitarian attitude implausible and believe it may even, in a way, reverse the appropriate order.

observe any thing but the perception."[12] I think that when we enter most intimately into what we call groups, we always stumble on the centers of consciousness of some particular element or other of the group. We can never catch the group's consciousness itself. Whatever we might think of Hume's theory of the self, I would favor something like a "bundle" theory of the group's consciousness. Even John Searle, himself a friend of a nonreductive conception of collective *intentionality,* disdains the idea of a "collective consciousness" as akin to that of a "super mind, floating over individual minds" and as an unattractive alternative to his own view. According to me, the sort of collective intentionality there can be in the absence of such a collective consciousness permits one sort of normative evaluation, but not another. The group can do good or bad things, even intentionally; and there can even be a derivative sense in which the group can *be* a good thing (in something like the sense that a machine gun is not as good a thing as a park bench). But without a unified locus of consciousness, the group's will cannot express its goodness in the way that the will of persons can express their goodness.

III. Normative Status Groups Can Have

It will be useful to contrast this view now with a line of thought that has been very nicely developed in essays by Philip Pettit and David Copp. Pettit writes: "There is a type of organization found in certain collectives that makes them into subjects in their own right,"[13] and Copp holds that "at least some collectives are agents, capable of acting intentionally."[14] As I have already indicated, I agree with these claims and, indeed, with most of the claims made by Pettit and Copp on behalf of their position. The view I would defend is consistent with their view; it represents an optional additional claim.

Pettit introduces his view by taking up the so-called Doctrinal Paradox identified by Lewis Kornhauser and Lawrence Sager.[15] This phenomenon can arise whenever a group is making a joint decision—by, for example, majority vote (other procedures are possible)—on the basis of several factors. The group could judge of each of the factors that it favors a given decision and yet vote as a majority in favor of the opposite decision. In a classic case, a three-member panel of judges is charged with deciding the guilt or innocence of a criminal defendant. According to the relevant

[12] David Hume, *Treatise of Human Nature,* ed. L. A. Selby-Bigge, 2d ed., revised and edited by P. H. Nidditch (Oxford: Oxford University Press, 1978), book 1, part 4.

[13] Philip Pettit, "Groups with Minds of Their Own," in Frederick Schmitt, ed., *Socializing Metaphysics* (Lanham, MD: Rowman and Littlefield, 2003), 167–93.

[14] David Copp, "On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy'," *Midwest Studies in Philosophy* 30 (2006): 194–221.

[15] Lewis Kornhauser and Lawrence Sager, "Unpacking the Court," *Yale Law Journal* 96 (1986): 82–117; Lewis Kornhauser and Lawrence Sager, "The One and the Many: Adjudication in Collegial Courts," *California Law Review* 91 (1993): 1–51.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

statute, the panel should find the defendant guilty if and only if it finds that he has satisfied both of the required elements (*e1* and *e2*). The panel agrees to let a majority vote on any issue represent its collective judgment on that issue.

The first judge holds that the defendant has satisfied *e1* but not *e2*: he finds the defendant innocent. The second judge holds that the defendant has satisfied *e2* but not *e1*: he too finds the defendant innocent. The third judge holds that both elements have been satisfied and thus that the defendant is guilty. Accordingly, the panel, as a collective, judges that the defendant has committed both elements of the crime: their majority vote is in favor of his satisfying *e1* and of his satisfying *e2.* But they also find him innocent, by majority vote. The panel is thereby out of compliance with the statute, in spite of its members: as a panel, it was to find the defendant innocent if and only if it found that he failed to satisfy both elements of the crime. Each member of the panel has observed that constraint, individually, but the panel as a whole, which is the deciding body, has not. What should the panel do? Finding the defendant guilty can seem unfair: a majority of the members of the panel finds him innocent. But there is also something wrong with finding him innocent: it was the panel, not the individual judges, that was charged with finding him guilty if and only if *it* found that he had committed both elements. And they have, as a group, found him guilty of each element. The state might reasonably object to the verdict of not guilty.

Pettit shows how this phenomenon generalizes along various dimensions.[16] On his analysis, the example points to a very general dilemma faced by groups trying to make a judgment: "They may maximize responsiveness to individual views, running the risk of collectively endorsing inconsistent sets of propositions. Or they may impose the discipline of reason at the collective level, running the risk of collectively endorsing a conclusion that a majority of them—perhaps even

[16] Indeed, there has been a great deal of recent literature on related issues. For a few examples, see Franz Dietrich, "A Generalized Model of Judgment Aggregation," *Social Choice and Welfare* 28, no. 4 (2007): 529–65; Franz Dietrich and Christian List, "Arrow's Theorem in Judgment Aggregation," *Social Choice and Welfare* 29, no. 1 (2007): 19–33; Franz Dietrich and Christian List, "Judgment Aggregation by Quota Rules: Majority Voting Generalized," *Journal of Theoretical Politics* 19 (2007): 391–424; Peter Gärdenfors, "An Arrow-like Theorem for Voting with Logical Consequences," *Economics and Philosophy* 22, no. 2 (2006): 181–90; Martin van Hees, "The Limits of Epistemic Democracy," *Social Choice and Welfare* 28, no. 4 (2007): 649–66; Lewis Kornhauser and Lawrence Sager, "The Many as One: Integrity and Group Choice in Paradoxical Cases," *Philosophy and Public Affairs* 32 (2004): 249–76; Christian List, "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective," *Episteme* 2, no. 1 (2005): 25–38; Christian List, "The Probability of Inconsistencies in Complex Collective Decisions," *Social Choice and Welfare* 24, no. 1 (2005): 3–32; Christian List, "The Discursive Dilemma and Public Reason," *Ethics* 116 (2006): 362–402; Christian List and Philip Pettit, "On the Many as One," *Philosophy and Public Affairs* 33, no. 4 (2005): 377–90; and Marc Pauly and Martin van Hees, "Logical Constraints on Judgment Aggregation," *Journal of Philosophical Logic* 35, no. 6 (2006): 569–85.

all of them—individually reject."[17] According to Pettit, groups that adopt the second alternative "deserve recognition as intentional *and* personal subjects." Of course, on my view, this last conjunction is not as innocent as it might seem: not every *intentional* subject satisfies all the relevant conditions for *personhood.*

To his credit, Pettit is sensitive to the possibility of a distinction here. He asks explicitly, "What distinguishes personal from merely intentional subjects?"[18] But his answer is that unlike merely intentional subjects, which must (just) exhibit "rational unity," persons must be such that they can be held responsible for failures to exhibit that unity: "Rational unification is a project for which persons must be taken to assume responsibility."[19]

Pettit has no hesitation in arguing that groups are, accordingly, institutional persons. Groups can be taken to assume responsibility for rational unification. But I think there is an important factor that he overlooks: by his own admission, groups can, or can fail to, "impose collective rationality." If they "bind themselves to the discipline of reason at the collective level," then they are what he calls social "integrates" and, accordingly, persons. Social groups that are not so bound contrast with such integrates. I think it is important, however, that there is no such option for individual persons. Persons are not such that they could have opted out of the rationality condition. Rationality is not a hypothetical imperative. Persons are essentially persons. Pettit's integrates are, at best, contingently so; but I think *a person* is something that nothing could contingently be.

Earlier in his discussion, Pettit suggests that groups subject to the sort of aggregative dilemmas to which he has pointed will be "under pressure" to impose the discipline of reason at the collective level—in effect, in the particular sort of example we are looking at, to let their collective judgment on what might be thought of as the "reasons" for a decision (e.g., *e1* and *e2*) dictate their decision, with no relevant independent vote on the conclusion itself (innocence). But the argument he offers for the existence of this pressure is of an entirely practical sort, involving the importance to the group of being "an effective and credible promoter of its assumed purpose."[20] Without the discipline of collective reason, the group will lose hold of its members and will lose respect with outsiders. Thus, Pettit argues, groups are under pressure to collectivize reason.[21] But this represents a striking contrast with the case of the individual person. The important thing about individual persons is not what practical pres-

[17] Pettit, "Groups with Minds of Their Own," 175.
[18] Ibid., 184.
[19] Ibid., 185.
[20] The practical character of this pressure is, on my view, an argument for an important difference (noted earlier) between an individual with inconsistent beliefs (of which he is aware) and a group with inconsistent beliefs.
[21] Pettit, "Groups with Minds of Their Own," 176–77.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

sure they are under to be rational: the demand on a person to be rational is *constitutive*.[22]

I think this supports a response to Pettit according to which even commitment to rational unification is insufficient for personhood. Just as intentional subjects can satisfy rational unity "whether or not they are aware of doing so" (to use one of Pettit's own phrases), so social integrates may be committed to rational unification without any awareness of being so committed—indeed, without any awareness at all. Pettit's mention of awareness reflects, I think, an implicit sensitivity to the normative significance of consciousness. Even when the group binds itself to rational unity, it cannot be aware of doing so—not consciously aware. All of the collective's mental states are, in a sense, blind, because they are all exhibited without the character of consciousness.

A tempting objection to Harry Frankfurt's "second-order volition" account of freedom may be applicable by analogy here.[23] Given the prospect of desires (for example, the desire to smoke) that we want not to be "effective" (in Frankfurt's terminology), it can seem unlikely that satisfaction of a second-order desire—itself at risk of being one such—could be decisive with respect to freedom. Suppose that rational unity, which involves coherence in one's commitments, is insufficient for personhood because it is possible for one to lack awareness of the unity. It is not promising that successful *commitment to* such unity might nevertheless be sufficient, notwithstanding the possible lack of awareness of that commitment. And if we focus rather on *why* Pettit finds that additional commitment to be important, what it is supposed to be sufficient *for*—susceptibility to a distinctive sort of evaluation—then I have been suggesting that without a unified center of consciousness to be aware of the commitment, the entity is not susceptible to criticism in the right sort of way and thus is not a person.

My view may harmonize with a claim espoused by Kornhauser and Sager:[24] "In paradoxical cases a group can rationally order its attitudes . . . over the applicable reasons . . . or it can rationally order its attitudes over the applicable outcomes . . . ; but it cannot [necessarily] do both. In contrast, an idealized individual who had worked out all the kinks could rationally order her views over both her reasons and her bottom-line choices or outcomes."[25] Accordingly, "however appealing the basic idea of holding groups to account as though they were individuals may be, it cannot without more be the grounds for Pettit's insistence on the priority

[22] The thought, on Pettit's behalf, that it is constitutive of the group that is a person (having become one) to be rational—in distinction to the group that is not a person—raises a difficulty about the group's taking on, anew, an essential property.

[23] On Frankfurt's account, very roughly, one acts freely just in case the desires one wants to be effective are indeed one's effective desires. See Harry Frankfurt, "Freedom of the Will and the Concept of a Person," *Journal of Philosophy* 68, no. 1 (1971): 5–20.

[24] There are, I believe, significant residual differences.

[25] Kornhauser and Sager, "The Many as One," 252.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

of reasons."[26] At a minimum, my position is—with Kornhauser and Sager's—an instance of what we might call *distinctivism:* the view that the overall normative situation of individual persons is distinctive and cannot ultimately be reproduced perfectly by groups.

Copp adds a new argument to the case for the sort of position he shares with Pettit. After a compelling discussion of the intuitive basis for his view, and a review of Pettit's arguments, Copp sketches an example akin to that of the panel of judges discussed earlier: in this case, there is a tenure committee at a university, and the university's regulations stipulate that tenure committees should promote candidates for tenure if and only if they satisfy the university's standards in the three areas of research, teaching, and service. As before, a majority of the committee finds the candidate to have satisfied each of the requirements for the promotion. But the panel votes to deny tenure, nevertheless. No individual member has voted against tenure without judging that the candidate has been unsatisfactory in at least one area. Nevertheless, the committee as a whole has done so. According to Copp, the committee (and through it, the university) is now open to criticism: the decision is arguably irrational. But no individual person is open to criticism. So the committee is a normative entity in its own right, with "normative autonomy."

Again, I agree with much of what Copp says here. Though I do not find the charge of irrationality in particular to carry weight in criticism of the group, objecting on that basis would be an instance of what Copp calls a "normative challenge" to his example, and I rather agree with him that "we should expect there to be cases in which the normative properties of a collective differ from those of any member."[27] My point is instead that the sort of normative properties in question represent a distinctive subcategory of a broader class of normative properties. And I find the difference in varieties of normative properties to be of significant independent interest.

Take a group of people that crowd around a gate in an airline terminal. The group has clogged foot traffic and interfered with the airline's boarding process. Even in such mundane cases, I think, we can see the sort of normative autonomy to which Copp directs our attention. No one member of the group has the normative properties exhibited by the group as such. In these cases, we do not have even the sort of coordination and planning that, according to some, produces an important subclass of group action.[28] Nevertheless, the relevant sort of normative evaluation is in order, or so it seems to me. That group shouldn't have done that—the

[26] Ibid., 255.

[27] Copp, "On the Agency of Certain Collective Entities," 218.

[28] See Michael Bratman, "Shared Intention," *Ethics* 104 (1993): 97–113; Bratman, "Dynamics of Sociality," *Midwest Studies in Philosophy* 30 (2006): 1–15; Margaret Gilbert, "The Structure of the Social Atom: Joint Commitment as the Foundation of Human Social Behavior," in Schmitt, ed., *Socializing Metaphysics,* 39–64; and John Searle, "Collective Intentions and

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press

group did something it shouldn't have done. The claim, with respect to any individual member of the group, that they shouldn't have done what they did doesn't have quite the same force. This is not unlike the case where, with little or no forethought or planning, an individual does something wrong—inconsiderately bringing up an inappropriate topic of conversation, say. If the individual is not really aware of what he is doing, then although one variety of normative evaluation seems perfectly appropriate, we are inclined to withhold another variety of evaluation. We are inclined to offer the very lack of awareness as mitigating the badness of the action, or at least to resist the ordinary inference to the badness of the *person,* however bad we may continue to think the action itself. I would like us to draw the corresponding inference for the case of groups: although the group has done something bad (and no member of the group need individually share the normative properties of the group along *that* dimension), because the group is not conscious of what it does, it cannot be assessed along another dimension.

IV. Conclusion

To complete the argument promised at the beginning of the essay, then, there is no problem for Kantian theories about aggregating individual responsibility into collective responsibility—because there is, on such theories, no such thing as collective responsibility at all. Just as consequentialism is alone in facing problems about aggregating the value of downstream effects *of* actions, because Kantian views hold that effects are irrelevant to (that sort of) normative status, so consequentialism is alone in facing a problem about aggregating responsibility *for* actions, because Kantian theories are best understood as focused on the will of a single unified consciousness.

Let me close by filling out the picture a bit more (if also a bit more speculatively), for the sake of completeness. The claim of this essay is that Kantian deontological theories (under a mild development of the view) and consequentialist theories exhibit a similar contrast (against each other) with respect to upstream phenomena as they do with respect to downstream phenomena. Although those theories *can* be interpreted as competing theories of a single topic—ethics—they can also be seen as offering independent characterizations of varieties of ethical status. On one way of developing this latter view,[29] deontological theories are *person-centered* normative theories, characterizing primarily the ethical agent and concerned with the status of actions only derivatively. Consequentialist theories, by contrast, are best seen as *act-centered* normative theories, focused

Actions," in Philip Cohen, Jerry Morgan, and Martha Pollack, eds., *Intentions in Communication* (Cambridge, MA: MIT Press, 1990), 401–16.

[29] I pursue the development of such a view in "Standard Bearers" (unpublished manuscript).

on accounting for ethical actions and concerned with agents only inasmuch as they are the agents of the actions in question.

But however the distinction is developed, once we allow that the theories do not compete, we open up the possibility of distinguishing the applicability of the relevant normative concepts with respect to various phenomena. An immediate example is that of the value of the consequences of an action: consequentialism is, by definition, concerned with that; deontological theories ignore it. The suggestion here is that consequentialist theories can (though, of course, one or another theory could simply forgo this possibility—unwisely, in my opinion) smoothly accommodate the phenomena of group action and collective responsibility, but deontological theories will ignore them. The family of deontological normative concepts is inapplicable to entities without a will, and, as with (possible) victims of dissociative identity disorder, groups lack a will, in the relevant sense, since they are not constituted by a single unified consciousness.

Although it is not entirely clear, substantively, that these views *must* go together, I think it will be no surprise to find functionalist conceptions of mind combined with consequentialist views in ethics together with compatibilist positions on free will, all connected with a more or less physicalist metaphysics. Persons are constrained to be a variety of spatiotemporal entity. The opposing worldview is dualistic, deontological, and libertarian. Persons are conceived as fundamentally a distinctive kind of normative entity.[30] If it were not so natural to combine otherwise opposing positions in ethics with correspondingly "opposed" positions on issues in philosophy of mind, action theory, and metaphysics generally, then I, for one, would be more hopeful for a local resolution of the dispute, a resolution that found in favor of one position and against the other. But I think the depth and richness of the opposition—and, for that matter, its historical resilience—recommends a more Humean attitude: "from this circumstance alone, that a controversy has been long kept on foot, and remains still undecided, we may presume, that there is some ambiguity in the expression, and that the disputants affix different ideas to the terms employed in the controversy."[31] It would be ironic, in a way, if it were an insight from Hume that helped enable a Kantian, rationalist, dualistic, deontological, and libertarian worldview.

Philosophy, University of Texas at Austin

[30] "Person" is a "forensick term," in Locke's nice phrase; see John Locke, *An Essay Concerning Human Understanding*, ed. P. H. Nidditch (Oxford: Clarendon Press, 1975), book II, chapter 27, sec. 26. And the appropriation of *merit*—what being forensic amounts to—is, for Locke, "founded in a concern for happiness, the unavoidable concomitant of consciousness; that which is conscious of pleasure and pain" (ibid.).

[31] David Hume, *An Enquiry Concerning Human Understanding*, ed. L. A. Selby-Bigge, 3d ed., revised with notes by P. H. Nidditch (Oxford: Clarendon Press, 1975), sec. VIII.

https://doi.org/10.1017/S0265052509090098 Published online by Cambridge University Press