Journal of Ethics and Social Philosophy
VOL. 27, NO. 3 · MAY 2024
https://doi.org/10.26556/jesp.v27i3.2005

A THEORY OF COLLECTIVE VIRTUE

Matthew Baddorf and Noah McKay

We say things like "Enron's greed led to catastrophe for many investors" and "the Eighty-Second Airborne was very brave." This language suggests that we believe that groups are capable of virtues and vices. It is hard to know what to make of this idea, however. Humans have virtues and vices due in part to our mental capacities, but attributing mental capacities to groups can sound absurd, like invoking vaguely Hegelian group spirits who work over and above their human members.[1] Pressure to avoid such results can lead to summative views of virtues and vices: on such views, a group's virtue or vice is just a result of "summing up" the same trait in its members. (So perhaps the Eighty-Second Airborne was brave just because most of its members were.) Unfortunately, this safely reductive view suffers from counterexamples, such as Lahroodi's example of a church committee that is closed-minded due to social pressures despite being made up of open-minded individuals.[2]

If we want to understand how groups can have important traits such as greed, bravery, and open-mindedness, we need a credible view that avoids the implausibility of Hegelian group minds and the counterexamples to summativism. Or so we think, anyway; and even if some readers think those Scylla and Charybdis safer than we do, we hope they will agree that there is room to attempt a middle course. Here, we set out a view that does just that. "Imitationism" is a kind of nonreductive theory that explains how a virtue can be genuinely collective without requiring collective minds.[3] We articulate and defend

1 We should note that we are not actually accusing Hegel of holding this view; this is a caricature of Hegel, not an exegesis of his social philosophy. Also, not everyone agrees that collective minds are implausible. For a position friendly to group minds, see Theiner, "A Beginner's Guide to Group Minds." And for arguments that some popular positions in philosophy of mind imply the existence of groups minds, see Theiner and O'Connor, "Emergence of Group Cognition"; and Schwitzgebel, "If Materialism Is True, the United States Is Probably Conscious."

2 Lahroodi, "Collective Epistemic Virtues," 287.

3 Theories of collective virtue are theories about normative properties of collectives. These properties have been of interest to analytic philosophers since the Second World War. (Unsurprisingly, interest sprung up shortly after the discovery of particularly shocking widespread human rights abuses in group contexts; My Lai prompted an important early

our theory in section 1, explain how it accounts for some examples of collective virtue in section 2, and address two objections in sections 3 and 4. But first, a few comments about the rationale for and scope of our project are in order.

A key feature of our view is that it attempts to do justice to the intentional nature of collective virtue without making commitments to collective minds or collective phenomenal consciousness that many find implausible. The idea, in other words, is that we can not only deny collectives minds of their own, but we can preserve an insight at the heart of opposition to collective minds—the importance of intentional content.[4] And we can do this without having to deny the possibility of nonreductive collective virtue.

We will not consider all sorts of groups. The sort of groups we are going to discuss are groups we call "collectives."[5] Collectives are organized and structured groups that can survive the departure of at least some of their members.[6]

---

edited volume on the subject (discussed in French, *Individual and Collective Responsibility*, vii). Most work has focused on collective moral responsibility (and the free will required for it). For recent examples, see Copp, "On the Agency of Certain Collective Entities"; Haji, "On the Ultimate Responsibility of Collectives"; Braham and van Hees, "Responsibility Voids"; Hess, "Free Will of Corporations"; Dempsey, "Corporations and Non-Agential Moral Responsibility"; and Baddorf, "Phenomenal Consciousness, Collective Mentality, and Collective Moral Responsibility."

Attention to collective virtue lagged somewhat but has increased in the last two decades. Much work on collective virtue has been oriented around attempts to establish and account for nonreductive collective virtue. In our judgment, however, we do not yet have a fully satisfactory theory. For example, Lahroodi, "Collective Epistemic Virtues," considers and Fricker, "Can There Be Institutional Virtues?" develops accounts drawn from Gilbert's idea of joint commitments ("On Social Facts"), but (as we will argue in section 2) these cannot account for the wide variety of collective virtue cases. Beggs, "The Idea of Group Moral Virtue," gives a theory that (as we understand it) is consistent with our own but not as developed. Jones, "Numerous Ways to Be an Open-Minded Organization," gives helpful examples of the variety of collective virtues but does not attempt to provide a unifying theory explaining them. For more examples of contemporary work on collective virtues, see Anderson, "Epistemic Justice"; Ziv, "Institutional Virtue"; Gowri, "On Corporate Virtue"; Cordell, "Group Virtues"; and Diamantis, "The Law's Missing Account."

4 For an argument readily adaptable for use against collective minds built around the importance of phenomenal consciousness, see Horgan and Kriegel, "Phenomenal Intentionality Meets the Extended Mind." One of the authors discusses the issue of collective minds further in Baddorf, "Phenomenal Consciousness, Collective Mentality, and Collective Moral Responsibility."

5 Sometimes these are called "corporate agents" (e.g., in the title of Björnsson and Hess's "Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents"). We will stick with the word "collective" and refer to virtues held by collectives as "collective virtues."

6 The metaphysics of collectives is a very difficult topic, and we do not intend to make any contribution to it here. For readers who would like some idea of the metaphysical framework we are assuming, though, we can say this: we suspect that collectives are artifacts

They are typically capable of group goals and actions, at least in some sense; and these are not necessarily shared by all their members.[7] The two of us going to the store together are not a collective; British Petroleum and the French government are. Collectives need not all be legally recognized or subject to a written code of rules, as these groups are. An informal chess club, organized by volunteers who rely only on informal social norms to govern their behavior, could well be a collective also.

There may not be any bright, clear line between the groups that are collectives and those that are not. Some cases may be hard to classify. The important thing is that our claims are only intended to apply to groups that are clearly collectives; whether they apply to other groups is a question we will not address.

We are interested in both virtues and vices, but here we will often just speak about virtues. We suspect most such statements will apply to vices as well.

With all that said, we can turn to motivating our view. Some readers may wonder why they should care about what is basically a classificatory debate. Almost everyone agrees that collectives can have dispositions that are similar to human virtues but not exactly like them: there is something to the claim that the Red Cross is compassionate, even if it is not compassionate in just the way a person might be. Why not say, as Cordell suggests, that what appear to be genuine collective virtues are really not virtues but just highly desirable structural features?[8]

We are somewhat sympathetic to this worry. But there are practical advantages to recognizing the existence of collective virtues and vices. First, and likely most importantly, recognizing collective virtue allows for a sort of broadly moral criticism of collectives that recognizing merely beneficial structural features does not. "Moral blame" is a complex and disparate set of phenomena, and we think that there is at least one sort of moral blame for which collectives are not apt targets.[9] But there are some practices that have gone by the title "moral blame" for which collectives can be apt targets.

---

that have humans as proper parts (and which are capable of persisting through changes to their human parts). See Baker, *The Metaphysics of Everyday Life*, for an influential discussion about the reality of artifacts, and Uzquiano, "Supreme Court and Supreme Court Justices," for an adaptation of Baker's account to groups. For some good recent discussion of the metaphysics of groups in general, see Ritchie, "What Are Groups?" and Epstein, "Ontological Individualism Reconsidered."

7 There has been a good deal of discussion recently about the nature of collective decision-making. For an authoritative summary of the ways that collective decision-making is importantly irreducible to that of individuals, see List and Pettit, *Group Agency*.

8 Cordell "Group Virtues."

9 One of the authors defends this claim in Baddorf, "Phenomenal Consciousness, Collective Mentality, and Collective Moral Responsibility."

Some of these practices are associated with attitudes such as disdain and admiration.[10] Paradigmatic examples of things we disdain and admire are human agents, but we can sensibly adopt these attitudes toward anything capable of a bad or good character. It is plausible that we can sensibly adopt these attitudes toward collectives; one can disdain a charity for its tendency for mission creep or admire a government for its justice. Collective virtues help account for and legitimize such attitudes. But mere structural or individual features—considered in themselves rather than as parts of the virtues they help constitute—do not.

Relatedly, consider the motivational effects of a belief that one's organization has a virtue or vice versus the belief that it merely has a feature with positive or negative effects. While the latter can certainly motivate one to preserve a collective's good features or reform its bad ones, we doubt that mere features have the motivational oomph most of us would get from learning that we are part of a virtuous or vicious organization. Collectives can have all sorts of positive and negative effects due to structural and individual features, many of which we do not have any strong obligation to encourage or discourage. But learning that your collective has a virtue involves learning more than this: it involves learning that you are participating in a morally praiseworthy endeavor. Recent empirical research suggests that this has a powerful effect on employees' willingness to identify with their employers.[11] Similarly, there is a natural horror one can feel when one discovers that one is participating in a morally diabolical group that is not felt simply by participating in a system with some equally undesirable consequences.

Finally, thinking in terms of collective virtues can also be practically advantageous for epistemic reasons. More particularly, it can help us conceptualize what it is that we are trying to get our collectives to do and be. Consider a scientific research group whose primary output is journal articles and that is thus analogous to an individual researcher. Thinking about their group in terms of the collective epistemic virtues that it does or does not exemplify can help the members of the team understand what it is that they are trying to make their group like. Making their group more virtuous might require moving beyond merely thinking about collective virtue, but that does not mean that it is not helpful to start with collective virtue. Skeptics might reply that one could, instead of starting with thoughts about what collective virtues the group should instantiate, start by thinking about what desirable structural features the group should have. But structural features combine with individual features in order to have their effects, and there are many different combinations of the

10 Here we draw from Shoemaker, *Responsibility from the Margins*, particularly ch. 1.

11 Chun, "Organizational Virtue and Performance."

two sorts of features that might result in a given effect. Thinking in terms of collective virtues would allow the group's members to start in the place that makes the most sense, with the sort of characteristic traits they want the group to have, rather than from the building blocks of those traits. Thinking in terms of collective virtues, then, can help us conceptualize what sort of agency-imitating collectives we should be striving for.

So much for the rationale. Now for the theory.

1. IMITATIONISM: A THEORY OF COLLECTIVE VIRTUE

Here is a summary of our view:

> *Collective Virtue* (CV): Collectives can possess virtues and vices; they can do so because they can possess reasonably broad and stable dispositions that functionally and intentionally imitate individual virtues and vices.

The gist of CV is just this: collectives have virtues because they have dispositions that meet certain conditions. Just what those conditions are needs explication. The basic idea, though, is that these collective dispositions resemble individual virtues closely enough to be virtues themselves.

The first notion in CV that needs some explanation is that of "reasonably broad and stable dispositions." By "stable dispositions," we mean dispositions that tend to endure over time; if a collective has one at a given time, then the collective will tend to have it in the future as well. By "broad," we mean that these dispositions trigger in an appropriately wide range of circumstances; a generous collective, like a generous individual, will usually not be arbitrarily generous in one situation but fail to be so in a similar one. Both stability and broadness admit of degrees; the claim that these dispositions are "reasonably" stable and broad is meant to indicate that collectives may not need a very high degree of either for their dispositions to count as virtues. Just as individuals need not always behave in accordance with a virtue in order to have it, so it is with collectives. Where the line ought to be drawn between virtue possession and lack is highly dependent on the virtue and is, of course, almost always difficult to discern precisely.[12]

12 The extent to which individuals have dispositions that meet these conditions has been the subject of a great deal of dispute in (individual) virtue theory over the past fifteen years. See Phillips, "Towards an Empirically Adequate Virtue Ethics," for discussion. We will assume that individuals have sufficiently broad and stable dispositions that they can have virtues. It is worth noting, though, that almost all parties to the debate over individual virtues think that individuals' dispositions are often less broad or stable than we tend to pretheoretically think. This gives us reason to think that collectives' dispositions may not need to be as broad or stable as we might pretheoretically think in order to be on as good a footing to qualify as virtues as those of individuals.

The second notion that bears comment in CV is that of functional imitation. What we mean here is that collective dispositions can play either the same or very similar functional roles as those played by individual virtues. For example, collective bravery might enable an organization to correctly determine what should be done in cases of danger and to implement its decisions; this role is identical to or similar to the role that bravery plays in the case of the brave individual.

The third notion to discuss—and the one that we suspect will provoke the most controversy—is that of intentional imitation. The idea is similar to that of functional imitation: collective dispositions can involve the same or very similar intentional content as those involved in individual virtues. Let us first say a bit more about what this means and why it matters, and then explain why we should think it is true.

What would it mean for a collective disposition to involve intentional content? To answer this question, we should first figure out what it is for an individual virtue to involve intentional content. The phrase "intentional content" is not often discussed with regard to virtues, but there is a straightforward sense in which they can involve intentional content: having a virtue can sometimes involve possessing intentional states. For example, courage involves a disposition to certain attitudes toward danger. These attitudes are intentional: they are directed toward danger. Some collective dispositions can involve similar states that are directed toward danger in the same way (or similar ways) as the individual attitudes involved in courage.

These collective states are not necessarily the same states as those in individuals. For example, individual courage likely involves certain beliefs and desires. We doubt that collectives have beliefs and desires, strictly speaking. Why not? Briefly: because collectives lack minds, and mental states, beliefs, and desires can only be had by creatures with minds. (If we are wrong about this, of course, the similarity with individual states is only strengthened.) But even if they cannot have beliefs and desires, collectives can have states that have the same intentional contents as beliefs and desires. Intentional content is a more plausible feature of collectives than mental states since there are many things with intentional content but no mental states: spoken and written communication and much art, for instance. (See section 4 for more discussion of this sort of intentional content.)

For example, the Biden administration has taken an official stance on immigration, and this official stance might (and probably does) have the same intentional content as a set of beliefs that could be held by an individual. When we say that collective dispositions can intentionally imitate individual virtues, we mean that they can involve intentional states with the same (or very similar) intentional content. We can call this the Content Imitation view.

Why is this important to Imitationism? In short, because some virtues are essentially intentional (or so we think). Courage is not merely a matter of behaving fearlessly in the presence of danger—it involves a certain attitude *toward* danger. Similarly, compassion is not merely a matter of behaving a certain way, but also of being concerned *about* the well-being of others. So, in order to be courageous or compassionate, collectives must be capable of imitating these attitudes and concerns. It would be exceedingly odd to admire a collective for its courage and compassion while denying that it had any concern for the vulnerable or stalwartness toward danger.

Why think that the Content Imitation view is true? Briefly, collectives can be fruitfully understood from the intentional stance. In other words, attributing intentional contents to states of collectives allows us to make sense of their behavior in the same way that attributing intentional contents to states of individuals does. That is why we often explain and predict the behavior of collectives in intentional terms: the sentence "the Roman Catholic Church opposes birth control" works as an explanation of some behaviors of the Roman Catholic Church because the Roman Catholic Church is related to birth control in something very like the way that an individual with a con-attitude toward birth control is.

Some philosophers have suggested that intentionality *consists in* intelligibility from the intentional stance.[13] If that is right, then collectives can certainly have intentional states. However, we note that the Content Imitation view does not require anything this strong. Even philosophers who deny that intelligibility from the intentional stance is sufficient for intentionality normally think that it counts as good evidence for intentionality. So, while we cannot be certain that collective states are sometimes intentional, we are reasonably confident that they are. We are not alone in this respect; many philosophers have defended detailed accounts of collective intentionality in recent years.[14] We do not endorse any one of these accounts in particular, but we think they render the Content Imitation view plausible. (We realize this brief explanation will not satisfy everyone—for reasons of scope, we cannot expand on it much. But in section 4, we try to anticipate an objection to this part of the account without moving too far afield.)

13 Dennett, in *The Intentional Stance,* holds that this is true for intentionality of all kinds. Kriegel has argued that it applies to derivative intentional content, even if nonderivative content is essentially linked to consciousness (*Sources of Intentionality,* ch. 4). See Tollefsen, "Organizations as True Believers," for an application of Dennett's view to collectives, and our section 4 for more on how our view is consistent with views like Kriegel's.

14 See Schweikard and Schmid, "Collective Intentionality," for an excellent survey, and Toumela, *Social Ontology*, for a developed and illuminating account.

To bring what has been said thus far together: we are saying that the fact that collectives can possess broad, stable traits that functionally and intentionally imitate human virtues is the reason that collectives can possess virtues and vices. The idea is that in cases of collective virtue, the traits we have described form an explanation for the collective virtue in question. We are not committed to a particular view of what sort of explanation is involved here; presumably it is not a causal one, but some other sort—probably some kind of constitution or grounding.[15]

Here is the thought behind this explanatory claim (however its details are worked out): we can determine whether something counts as a virtue or vice by comparing it with paradigmatic instances of such things. (Paradigmatic instances of a virtue might include generosity, or courage, or other examples of traits commonly attributed to individuals.) Collective virtues and vices are not paradigmatic instances, but they are sufficiently similar to paradigmatic instances to count as instances. In large part, this is because they are, like paradigmatic instances, broad and stable, and because they both play the same or similar functional roles and involve the same or similar intentional content as those of individual virtues and vices.

We are not saying that collectives can possess all virtues and vices; there may be some virtues that involve things that collectives lack. In fact, when we attribute some trait to a collective—say, open-mindedness—it may be that the word "open-mindedness" refers to a somewhat different trait than when we use the word with respect to an individual. Obviously, the two traits could not be too dissimilar, or else it would not make sense to use the same word for each, but it is nonetheless possible that the collective trait is different from the individual one. This is worth bearing in mind when we consider a virtue like courage, which we might be tempted to think that collectives cannot possess. (Maybe it involves rightly facing a phenomenally conscious fear.) Even if they cannot possess human courage, it is natural to think that they can possess some collective analogue of it, and that is all our theory needs. Indeed, CV could be true even if collective virtues are not close analogues of any individual virtue. So long as the collective traits count as virtues due to their "formal" similarity to paradigmatic cases of individual virtue as described by CV, Imitationism allows for the possibility of "alien" collective virtues that are little like the virtues of human beings.

One last point about CV: one might wonder why we are not simply functionalists about virtue, claiming that possessing a virtue is a matter of instantiating

15 Note that while we are providing explanations, we are not providing an analysis of the concept of collective virtue. For some reasons why we do not think this is a helpful endeavor, see Huemer, "Failure of Analysis and the Nature of Concepts."

certain functional roles.[16] If we were, then collective virtue would be comparatively straightforward, perhaps even easy to establish, and certainly the concept of a functional role has come up often enough thus far that it is plausible to think that functional roles are important for understanding virtue. In part, we do not endorse this view because it is not at all clear to us that the virtues possessed by individual human beings are actually simply a matter of functional roles. Phenomenal consciousness may well play an ineliminable part of certain human virtues. (In addition to the example of courage above, consider virtues that involve empathy, such as compassion.) Or perhaps not: but we think neutrality about this issue is an advantage of our view. If Imitationism is correct, then the existence of collective virtue is not hostage to the fortunes of functionalism.

2. Examples of Collective Virtue

Soon we will consider how our view can deal with some objections, but first, we want to consider some further examples of collective virtue and vice to show how our view can account for them. Let us start with a simple case of collective virtue, one similar to a case suggested by Donald Beggs.[17] Suppose a quilting group consisting of white, middle-class Americans regularly sells their quilts and donates the proceeds to charities. The members of the group each think that the group should not donate funds to charities that simply benefit people like themselves (e.g., charities that focus on finding cures for diseases that threaten the affluent); instead, they want to donate to charities that help the most disadvantaged members of society. So, each member is careful to ensure that the group's chosen charities fit with this goal. In particular, they exert pressure on the committee that selects charities to choose appropriately. In this case, the group has a virtue: the group cares for those who are most disadvantaged, despite the fact that they are often socially, ethnically, and economically different from the group's members. Beggs calls this virtue "radical tolerance," but we could just as aptly name it "compassion."[18]

16 We would like to thank an anonymous reviewer for pressing this question.

17 This case is a variation on Beggs's example of a quilting group ("The Idea of Group Moral Virtue," 467–68). In his original case, the members do not each exemplify the care and concern mentioned below.

18 It is worth remembering here that, on our view, the trait of collective compassion might not be identical to the individual trait that gives it its name. Individual compassion might require, say, a phenomenally felt emotion, and probably collectives cannot have those. See the previous section for more discussion (and thanks to an anonymous reviewer for helping us see that a reminder is in order here).

Now, the individual members of this group have the same virtue as the group. So, we do not give this example as an example of a group whose virtue is in some strong way irreducible to those of its members. It does show a simple and straightforward case of collective virtue, however, and one that is aptly explained by CV. For in this case, the collective has traits that meet CV's requirements. It is disposed to perform actions that are functionally and intentionally similar to an individual's virtuous choice: the group selects charities *because* donating to those charities is likely to improve the lot of those who are worst off. And the vigilance of the members ensures that these dispositions are reasonably broad and stable; in a variety of different sorts of fundraising cases, the quilting group will make these choices, and it is likely to retain this disposition for some time. We think it is plausible that these sorts of considerations explain the fact that the quilting group has the virtue of compassion. It is by way of these traits that it does so. In this case, the quilting group has the appropriate dispositions because each of the individual members has the same virtue. But the virtues of individuals do not automatically "rise" to the level of the group; each individual quilter could be compassionate and simply express it outside the group. So, CV is playing some explanatory role here, despite the fact that the example is not an obvious counterexample to summativism.

Let us turn to a less apparently reductive case. This next example involves a collective vice—namely, collective carelessness. It is based on an example given by Kendy Hess, who gives it as part of an argument for irreducibly collective moral responsibility, but it is equally interesting as an example of a collective with a vice not shared by its members. In Hess's example, ACME corporation has employees who are committed to protecting the environment.[19] However, after ACME decides to produce steel additives, the distributed nature of its decision-making about the details causes a problem. In Hess's account:

> Member *A* requests proposals from Departments α, β, and γ,
> Member *B* picks the one from α and modifies it slightly to reduce costs,
> Member *C* modifies the proposal to improve materials handling,
> *D* modifies the proposal to improve efficiency,
> *E* modifies it to improve health and safety compliance (less worker exposure),
> *F* modifies it to use different (nationally available) chemicals, and
> *G* modifies it to reduce costs again.
>
> In the end, as a result of these piecemeal modifications and others during implementation—each innocuous and rational enough within

19 Hess, "Free Will of Corporations (and other Collectives)," 247–48.

> its own limited sphere—the new production line results in a continuing discharge that pollutes a local river.[20]

We can suppose that the pollution involved is of a sort and level that is inconsistent with appropriate regard for the environment. So, it seems that ACME corporation is environmentally careless. Although the individuals involved may not have been able to realize it (given their limited individual knowledge and time), the fact that their actions together resulted in the discharge shows that ACME's decisions do not take environmental impacts into account. It is, therefore, plausible that ACME has a vice of disregard for the environment insofar as its method of decision-making does not seem to provide for environmental effects, at least when they are not obvious.

One might think that ACME lacks a disposition stable enough to count as a vice since, for all we know, given Hess's description, a slight change in circumstances could have resulted in a very different outcome.[21] Maybe if *G*, for instance, had not made that last modification, no pollution would have occurred. But note two things here: first, we could foreclose this possibility by assuming that the environment is extremely sensitive to ACME's activities and that the few ways of carrying out those activities in an environmentally responsible way are less efficient, affordable, and safe for personnel than the alternatives. On these assumptions, ACME is likely to cause some environmental damage in any given case. Second, and more importantly, vices of carelessness do not need to have a high probability of resulting in bad outcomes to be broad and stable enough to count as vices. Most careless drivers do not cause crashes on any given trip, and when they do cause a crash, it is often the case that the crash might easily have been avoided if they had made some apparently unimportant change (such as driving down a different road). Similarly, pollution need not be a high likelihood event for ACME to have a vice due to its lack of care to ensure such pollution does not occur.

To sum up: in this case, the corporation has dispositions to make decisions in certain ways, and these dispositions are functionally and intentionally similar to an individual who behaves with careless disregard for the environmental consequences of their actions. CV thus captures ACME's vice nicely.

Unlike the previous example, this is a case where the individuals involved do not share the collective trait: all the employees of ACME (and all the owners as well) might be environmentally virtuous yet not realize that their collective efforts make ACME vicious. We could also imagine a reversed case, in which

20 Hess, "Free Will of Corporations (and other Collectives)," 248.

21 We would like to thank an audience at the University of Rochester for making this objection.

most (perhaps all) employees were not environmentally virtuous, but a carefully followed company policy of checking for and correcting environmental problems resulted in a virtuous company. Either way, this sort of case is not susceptible to a reduction to the virtues and vices of individuals.[22] The reversed case is also an example of collective virtue that cannot be explained by theories that rely on the notion of a joint commitment to collective virtue, such as those of Fricker and Lahroodi, since in the reversed case, the employees of ACME never jointly commit themselves to caring for the environment.[23] These sorts of cases, by contrast, are nicely explained by CV.

The last case we want to look at now illustrates how Imitationism can account for virtues in "invisible hand" cases: cases where a collective outcome is the result of conflicting individual behavior. Miranda Fricker considers a case like the following one.[24] A jury is made up of biased individuals, but because their biases cancel each other out, the jury reaches a fair verdict. Fricker acknowledges that there is a sense in which the jury is fair-minded, considering all the evidence of the case and weighing it appropriately before coming to the most reasonable conclusion. Yet she thinks the jury unvirtuous because the members are all biased—the jury is fair-minded only because their biases cancel each other out. (Presumably, when juror *A* ignores some key bit of evidence because it does not conform to her bias, juror *B* points out the importance of the evidence, and the group ends up weighing it appropriately.) In this case, Fricker thinks there is no collective virtue because the resulting correct verdict was not the result of any good motivation or skill—it is merely the chance result of the individuals who happened to be selected.

We think that Fricker may be right to think that collectives whose good actions are the result of happenstance are not thereby virtuous. And CV can account for this since these good results are not the product of broad and stable dispositions. However, it is also possible that the jury's fair-minded result is not mere happenstance. (Fricker does not discuss this possibility.) Suppose the jury was formed by a legal system that reliably forms juries with individuals of sufficiently varied backgrounds and biases that juror's biases are likely to cancel each other out. This might or might not be the product of intentional design of the legal system. (We have no position on the extent to which any actual

22 We think that at least some cases of collective virtue and vice are nonreductive in a stronger way: they do not supervene on (or reduce to) any set of intrinsic properties of individual agents. But defending that claim would take us too far afield here.

23 Fricker, "Can There Be Institutional Virtues?"; and Lahroodi, "Collective Epistemic Virtues."

24 See Fricker, "Can There Be Institutional Virtues?" 239. Fricker's original case concerned a debating society, but the point is the same.

legal system manages to achieve this, but it is our impression that the United States' adversarial jury selection system may be designed to produce this sort of effect.) In this case, it is not a coincidence that the jury behaves fair-mindedly. The jurists have been placed together in such a way that their behavior will lead to an outcome far better than the outcome that any individual jurist would have produced. This seems to us to be a case of a collective virtue. Our theory would handle this case by saying that the jury's reliable behavior is a collective functional and intentional analogue of an individual's disposition to be fair-minded; thus, the jury's finding would be interpreted as a result of a collective virtue.[25]

This case illustrates an advantage of Imitationism. We want to be able to differentiate between cases where a process simply happens to result in a good outcome and cases where there is a genuine collective virtue. Existing joint commitment accounts of collective virtue do this by requiring a joint commitment among individuals to the virtue, but this leaves out cases of collective virtue such as this one. CV, by specifying that collective virtues must be reasonably broad and stable, can do the same work and account for these cases.

In sum: we have good reason to believe that collectives can have virtues through broad and stable functional and intentional imitation of characteristics of paradigmatic individual virtues. This view allows collective virtues to be nonreductively held in a reasonably strong sense, and accounts nicely for a range of different sorts of collective virtue.

3. An Objection: Truly Agential Virtues?

One major objection is likely to linger in the minds of some readers: Are the collective virtues we have described really virtues in the sense that most ethicists and epistemologists use the term? Typically, the sorts of virtues that get philosophical attention are what one might call robustly agential: they have to do with advanced capacities for decision-making—capacities that require advanced capacities for reflection (e.g., on the value of the virtue in question) that Imitationism does not capture. This sort of objection has been developed by Sean Cordell into a dilemma: either a theory of collective virtue claims

25 Would it be a collective virtue of the jury or of some other collective, such as the legal system? We are not sure. One might think that the jury could not have the sort of broad disposition necessary. For juries might have all their members essentially (replacing a single member might result in a new jury), and if the trial were different, then different jurists might have been selected—and if that is true, perhaps the jury does not exist in enough different possible circumstances for it to be capable of a broad disposition to be fair-minded. But if so, we think that it would be reasonable to think that the collective virtue producing the fair-minded outcome in this case could be a virtue on the part of the legal system as a whole: a virtue of selecting good juries.

that advanced psychological capacities are required for collective virtue, or it does not.[26] If it does, then it is hard to see how collectives can possibly meet the requirement (at least if one is trying hard to avoid Hegelianism or group minds, as we have). If it does not, then it is not clear that collective virtues are, after all, similar enough to the sort of moral and epistemic individual virtues that matter to be worthy of the virtue label. Cordell concludes that we should replace the idea of collective virtue with a somewhat more reductive account of desirable and undesirable structural features of collectives.

Both horns, we think, can be resisted. Consider Cordell's second horn, on which collectivists like us claim that collectives can have virtues *without* having advanced psychological capacities. Suppose that we humans need these sorts of psychological capacities in order to have full agential virtues. Even so, it could be that collective virtues can still be similar enough to our own to count as virtues without such collective capacities. For the collective might still be capable of virtues that are partially agential. Imagine a spectrum: on one side, there are beneficial qualities without any agential features. (Knives, for example, can have beneficial qualities such as sharpness, though they lack agential virtues.) On the other side of the spectrum are fully agential human virtues, virtues that involve the capacity for reflection on value. Even if they are not at the same end of the spectrum with fully human virtues, collectives might have virtues that are considerably closer to that end than to merely beneficial qualities. Collectives (unlike knives) could still be in states that are functionally and intentionally very similar to those of individuals with beliefs and desires. This imitation can result in robustly nonreductive cases of collective virtue and vice in ways that mimic those of individual agents. It is plausible that this is enough to make collective traits agential in ways that make them count as genuine virtues, in part because they are apt for the same kinds of appraisal and emotional response as virtues (as when we are, for instance, ashamed of associating with a vicious collective). In other words, they could be agential enough even if they are not as agential as us.

Still, we hold out hope that Cordell's challenge can be met by addressing the first horn of his dilemma.

On the first horn of Cordell's dilemma, virtues require that their possessor have certain psychological capacities: capacities for reflection, or at least the ability to act upon the value of those virtues. We are not sure whether such capacities are required for virtue in individual human agents, but suppose that they are. On our view, collectives might well be capable of such capacities—or

26 See Cordell, "Group Virtues," 53–56. Cordell's original target was Donald Beggs's account (see Beggs, "Idea of Group Moral Virtue," 464–66). We are adapting and generalizing Cordell's dilemma somewhat.

of capacities that are functionally and intentionally similar enough to their individual counterparts to suffice for genuinely agential virtue. Key to any such account would be to work out functional and intentional equivalents of the required capacities that are equivalent in that they fulfill the reason(s) such capacities are required. What these reasons are depends on what the right account of individual virtue says about why these psychological capacities are required in individual cases.[27]

We will not pursue that project here, but we will note one promising avenue. It seems plausible that some individuals could reflect on the value of a collective virtue in such a way that their reflection instantiated collective reflection. (The collective would be functionally and intentionally imitating such reflection by way of the individual reflection.)[28] This would not make the collective redundant; the virtue would be the collective's since (by hypothesis) other elements of the collective (individuals as well as nonhuman things like computer systems and bylaws) will be needed to instantiate other aspects of the collective virtue. If all this is correct, then Imitationism need not make the concession that collective virtues are not as fully agential as individual human virtues. Collectives can, in their own way, have virtues as robustly agential as those of individuals.

4. Another Objection: Genuine Intentional States?

In section 1, we claimed that collectives are able to imitate human virtues in part by having states with the same intentional contents as states that partly constitute human virtues. We think this is reasonable since collectives are sometimes best understood from the intentional stance, and this is good evidence that some of their states are intentional. However, some philosophers might worry that the Content Imitation view is incongruent with a growing research project in the philosophy of mind, which Kriegel dubs the "phenomenal intentionality research program," or "PIRP."[29] PIRP theorists hold that there is some kind of tight connection between intentional states and phenomenal consciousness.

27 For promising examples of how this sort of argument can go in a discussion of the psychological requirements for collective moral responsibility, see Björnsson and Hess, "Corporate Crocodile Tears?"; and Collins, "I, Volkswagen."

28 An anonymous reviewer has suggested that this proposal might allow a collective to have a mind after all. There is a lot that could be said here, but in brief, while this might be true, we suspect that even here there would not be a collective mind, but simply a collective making use of an individual mind. To put it another way, the collective would have no mind of its own, but would merely appropriate the reflection of a human mind. More inspiration for how this might work could be gotten from Collins, "I, Volkswagen."

29 Kriegel, *Sources of Intentionality*.

Since most people, including us, think collectives cannot be phenomenally conscious, it is difficult to see how they could be in states with intentional content if some version of PIRP is true.

The nature of intentionality is among the most vexed topics in all of philosophy, and we will not try to add to the debate about it here. We will simply offer a few reasons to think that, even if one of the aforementioned theories of intentionality is true, this does not *prima facie* pose a serious problem for Imitationism. In fact, we think our view is consonant with PIRP in that we, like PIRP theorists, are keen to avoid the overbroad ascriptions of mental states that can come from excessive reliance on functional analyses.

Even those sympathetic to PIRP typically allow that some nonconscious states can have intentional content in a derivative way.[30] For example, most would allow that a written token of the sentence "All men are mortal" has the same intentional content as the thought that all men are mortal, though only derivatively. Similarly, many allow that subconscious mental states have intentional content derivatively, and some of these subconscious states are partly constitutive of virtues and vices. A subconscious bias against members of a certain racial group, for instance, might have that racial group as its intentional object, and this might be part of what makes the bias genuinely vicious.

So, there is reason to think that some nonconscious states can have derivative intentional content, even if phenomenal consciousness is necessary for underived content, and some such states can apparently partly ground or constitute virtues and vices. To those readers who are allies of PIRP, we submit that the intentional states with which the Content Imitation view is concerned have their content derivatively.[31] As we said before, we do not have a theory of collective or derivative intentionality. But we hope that the right theories will accommodate collective intentional states like the ones Imitationism requires.

30 Not all PIRP theorists believe this. For strong eliminativist theories of nonconscious intentional content, see Strawson, "Intentionality and Experience" and "Real Intentionality"; and Georgalis, *Primacy of the Subjective*. We think it strains credulity to insist that there are no subconscious intentional states or that sentence tokens have no intentional content.

31 For various accounts of derived intentional content among PIRP theorists, see Searle, *Rediscovery of Mind*, ch. 7; Loar, "Reference from the First-Person Perspective"; Horgan and Tienson, "Intentionality of Phenomenology and the Phenomenology of Intentionality"; Horgan and Graham, "Phenomenal Intentionality and Content Determinacy"; and Kriegel, *The Sources of Intentionality*, ch. 4. We think a satisfactory account will allow for collective intentionality. For example, Kriegel argues that a nonconscious state has derivative intentional content if an ideal observer approaching the world from the intentional stance would ascribe intentional content to that state. On a view like this, the Content Imitation thesis is unproblematic: collectives have derivative intentional states because an ideal observer would interpret some of their states as intentional.

Imitationism can, then, deliver robustly agential virtues even given theories of intentional content that accord great importance to phenomenal consciousness, all while eschewing Hegelian group minds.[32]

Walters State Community College
matthew.baddorf@ws.edu

Purdue University
mckay24@purdue.edu

REFERENCES

Anderson, Elizabeth. "Epistemic Justice as a Virtue of Social Institutions." *Social Epistemology* 26, no. 2 (2012): 163–73.

Baddorf, Matthew. "Phenomenal Consciousness, Collective Mentality, and Collective Moral Responsibility." *Philosophical Studies* 174, no. 11 (November 2017): 2769–86.

Baker, Lynne Rudder. *The Metaphysics of Everyday Life: An Essay in Practical Realism*. Cambridge: Cambridge University Press, 2007.

Beggs, Donald. "The Idea of Group Moral Virtue." *Journal of Social Philosophy* 34, no. 3 (September 2003): 457–74.

Björnsson, Gunnar, and Kendy Hess. "Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents." *Philosophy and Phenomenological Research* 94, no. 2 (March 2017): 273–98.

Braham, Matthew, and Martin van Hees. "Responsibility Voids." *Philosophical Quarterly* 61, no. 242 (January 2011): 6–15.

Chun, Rosa. "Organizational Virtue and Performance: An Empirical Study of Customers and Employees." *Journal of Business Ethics* 146, no. 4 (December 2017): 869–81.

Collins, Stephanie. "I, Volkswagen." *Philosophical Quarterly* 72, no. 2 (April 2022): 283–304.

Copp, David. "On the Agency of Certain Collective Entities: An Argument from 'Normative Autonomy.'" *Midwest Studies in Philosophy* 30, no. 1

32 We would like to thank Abby Brengle, Sean Cordell, Bill FitzPatrick, R. E. Fulton, Salim Furth, Kate Ritchie, Michael Robinson, Bill Rowley, Brett Sherman, and anonymous reviewers for the *Journal of Ethics and Social Philosophy* for their feedback and discussion. Similar thanks for constructive feedback are due to audiences at the 2016 Pacific APA, the University of Rochester, the 2017 Philosopher's Cocoon Conference, and the 2018 International Social Ontology Society.

(September 2006): 194–221.
Cordell, Sean. "Group Virtues: No Great Leap Forward with Collectivism." *Res Publica* 23, no. 1 (February 2017): 43–59.
Dempsey, James. "Corporations and Non-Agential Moral Responsibility." *Journal of Applied Philosophy* 30, no. 4 (November 2013): 334–50.
Dennet, Daniel C. *The Intentional Stance*. Cambridge, MA: MIT Press, 1989.
Diamantis, Mihailis E. "The Law's Missing Account of Corporate Character." *Georgetown Journal of Law and Public Policy* 17 (2019): 865–96.
Epstein, Brian. "Ontological Individualism Reconsidered." *Synthese* 166, no. 1 (January 2009): 187–213.
French, Peter A., ed. *Individual and Collective Responsibility*. Rochester, VT: Schenkman Books, 1998.
Fricker, Miranda. "Can There Be Institutional Virtues?" In *Oxford Studies in Epistemology*, vol. 3, edited by Tamar Gendler and John Hawthorne, 235–52. Oxford: Oxford University Press, 2010.
Gowri, Aditi. "On Corporate Virtue." *Journal of Business Ethics* 70, no. 4 (February 2007): 391–400.
Haji, Ish. "On the Ultimate Responsibility of Collectives." *Midwest Studies in Philosophy* 30, no. 1 (September 2006): 292–308.
Hess, Kendy M. "The Free Will of Corporations (and other Collectives)." *Philosophical Studies* 168, no. 1 (March 2014): 241–60.
Horgan, Terence, and George Graham. "Phenomenal Intentionality and Content Determinacy." In *Prospects for Meaning*, edited by Richard Schantz, 321–44. Amsterdam: de Gruyter, 2012.
Horgan, Terence, and Uriah Kriegel. "Phenomenal Intentionality Meets the Extended Mind." *Monist* 91, no. 2 (April 2008): 347–73.
Horgan, Terence, and John Tienson. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In *Philosophy of Mind: Classical and Contemporary Readings*, edited by David J. Chalmers, 520–33. New York: Oxford University Press, 2002.
Huemer, Michael. "The Failure of Analysis and the Nature of Concepts." In *The Palgrave Handbook of Philosophical Methods*, edited by Chris Daly, 51–76. New York: Palgrave MacMillan, 2015.
Jones, Todd. "Numerous Ways to Be an Open-Minded Organization: A Reply to Lahroodi." *Social Epistemology* 21, no. 4 (2007): 439–48.
Georgalis, Nicholas. *The Primacy of the Subjective: Foundations for a Unified Theory of Mind and Language*. Cambridge, MA: MIT Press, 2006.
Gilbert, Margaret. *On Social Facts*. Princeton: Princeton University Press, 1989.
Kriegel, Uriah. *The Sources of Intentionality*. New York: Oxford University Press, 2011.

Lahroodi, Reza. "Collective Epistemic Virtues." *Social Epistemology* 21, no. 3 (2007): 281–97.

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

Loar, Brian. "Reference from the First-Person Perspective." *Philosophical Issues* 6 (1995): 53–72.

Phillips, Kathryn. "Towards an Empirically Adequate Virtue Ethics." PhD diss. University of Rochester, 2015.

Ritchie, Katherine. "What Are Groups?" *Philosophical Studies* 166, no. 2 (November 2013): 257–72.

Schweikard, David P., and Hans Bernhard Schmid. "Collective Intentionality." *Stanford Encyclopedia of Philosophy* (Fall 2021). https://plato.stanford.edu/archives/fall2021/entries/collective-intentionality.

Schwitzgebel, Eric. "If Materialism Is True, the United States Is Probably Conscious." *Philosophical Studies* 172, no. 7 (July 2015): 1697–721.

Searle, John R. *The Rediscovery of Mind*. Cambridge, MA: MIT Press, 1992.

Shoemaker, David. *Responsibility from the Margins*. Oxford: Oxford University Press, 2015.

Strawson, Galen. "Intentionality and Experience: Terminological Preliminaries." In *Phenomenology and Philosophy of Mind*, edited by David Woodruff Smith and Amie L. Thomasson, 41–66. Oxford: Oxford University Press, 2005.

———. "Real Intentionality 3: Why Intentionality Entails Consciousness." In *Real Materialism: and Other Essays*, edited by Galen Strawson, 281–306. Oxford: Oxford University Press, 2008.

Theiner, Georg. "A Beginner's Guide to Group Minds." In *New Waves in Philosophy of Mind*, edited by Mark Sprevak and Jesper Kallestrup, 301–22. New York: Palgrave MacMillan, 2014.

Theiner, Georg, and Timothy O'Connor. "The Emergence of Group Cognition" In *Emergence in Science and Philosophy*, edited by Antonella Corradini and Timothy O'Connor, 78–117. Abingdon-on-Thames: Routledge, 2010.

Tollefsen, Deborah "Organizations as True Believers." *Journal of Social Philosophy* 33, no. 3 (Fall 2002): 395–410.

Toumela, Raimo. *Social Ontology: Collective Intentionality and Group Agents*. New York: Oxford University Press, 2013.

Uzquiano, Gabriel. "The Supreme Court and the Supreme Court Justices: A Metaphysical Puzzle." *Noûs* 38, no. 1 (March 2004): 135–53.

Ziv, Anita Konzelmann. "Institutional Virtue: How Consensus Matters." *Philosophical Studies* 161, no. 1 (October 2012): 87–96.