

# Phenomenal consciousness, collective mentality, and collective moral responsibility

Matthew Baddorf<sup>1</sup> 

© Springer Science+Business Media Dordrecht 2016

**Abstract** Are corporations and other complex groups ever morally responsible in ways that do not reduce to the moral responsibility of their members? Christian List, Phillip Pettit, Kendy Hess, and David Copp have recently defended the idea that they can be. For them, complex groups (sometimes called collectives) can be irreducibly morally responsible because they satisfy the conditions for morally responsible agency; and this view is made more plausible by the claim (made by Theiner) that collectives can have minds. In this paper I give a new argument that they are wrong. Drawing on recent work in the philosophy of mind (what Uriah Kriegel calls “the phenomenal intentionality research program”) and moral theory (David Shoemaker’s tripartite theory of moral responsibility), I argue that for something to have a mind, it must be phenomenally conscious, and that the fact that collectives lack phenomenal consciousness implies that they are incapable of accountability, an important form of moral responsibility.

**Keywords** Collective moral responsibility · Moral responsibility · Accountability · Intentionality · Phenomenal consciousness

---

✉ Matthew Baddorf  
mattbaddorf@gmail.com

<sup>1</sup> University of Rochester, 1670A Mt. Hope Ave, Rochester, NY 14620, USA

## 1 Introduction

Are corporations, or other similarly complex human groups (which we can call “collectives”) capable of the sort of moral responsibility had by individual human agents?<sup>1</sup> One sense in which someone might think they are is quite reductive: a collective<sup>2</sup> might be morally responsible just by virtue of having some of its decisions made by a morally responsible member. The collective would only be morally responsible in a derivative way, insofar as some morally responsible individual acted within some official capacity. Some philosophers, however, among them Kendy Hess, David Copp, Christian List, and Phillip Pettit, hold a stronger view [see e.g. (Hess 2010, 2013; Copp 2006; List and Pettit 2011)]. They maintain that some collectives can be morally responsible in ways that do not reduce to the moral responsibility of their members, and hence that the moral guilt which can attach to morally bad decisions can be held by collectives in ways which are not distributed amongst their members. Collectives, they believe, can be so functionally sophisticated that our best theories of agency count them as acting in morally responsible ways—with collective responsibility no more reducible to individual responsibility than your responsibility is reducible to “responsible” neurons.<sup>3</sup> Theoretically, at least, the group could be guilty even if all the members are innocent.<sup>4</sup> (Of course, usually individuals within groups are individually responsible

<sup>1</sup> Like most participants in this discussion, I am using “moral responsibility” to refer to what has sometimes been called “backwards-looking” moral responsibility—roughly, being able to be morally accountable for something one has done (or failed to do). I am not addressing the question of “forward-looking moral responsibility,” which is roughly the ability to be morally obligated to do/not do something. There are interesting questions about the relationship between moral responsibility and moral obligation, but I won’t pursue them here. And, as I note below, the sort of moral responsibility at issue here is really *non-reducible* moral responsibility: moral responsibility which is not had derivatively in some way by virtue of someone else’s moral responsibility. So, for example, one might think that a group can be morally responsible for some group action by virtue of the fact that its members are individually responsible for that action. In this case, the group would be morally responsible in a reductive way. I will set aside reducible moral responsibility for this discussion, but only use the adjective “irreducible” to modify the term “moral responsibility” when it seems important to remind the reader of the distinction. (Thanks to an anonymous reviewer for pointing out the need to say more about what I mean here.)

<sup>2</sup> By “collective” I mean any group that might plausibly be thought to be irreducibly morally responsible. Some use the term “corporate agent,” rather than “collective;” I don’t mean for anything to hang on my terminological choice.

<sup>3</sup> Hess (2013) provides an excellent defense of this notion. List and Pettit (2011) provide a defense and extensive discussion of the sorts of “deliberative dilemma” cases that can be used to make a kind of non-reductionism about collective actions plausible (see particularly chapters 1–3). And Copp (2006) also defends the idea that our understanding of agency extends agency to some groups; on page 220 he challenges those who disagree to provide a “solid rationale” for rejecting this view. This paper is an attempt to do so.

<sup>4</sup> Although rare, cases where a group might be guilty despite the fact that none of the individuals involved seem responsible exist. Here’s an example: In 1979, an Air New Zealand Limited flight from Auckland to McMurdo Station crashed into Mt. Erebus, killing everyone on board. The pilot would almost certainly have avoided the mountain if he had been informed that the plane’s autopilot route had been altered. Pilots were typically told about changes to their route orally, and informally, shortly before takeoff, but for some reason this was not done. Subsequent investigation found that no single individual had responsibility for informing the pilot. Furthermore, informal arrangements for notifying pilots were standard practice among airlines at the time. Investigators concluded that neither the employees on duty

for their own actions—no one in this discussion is trying to get individuals off the hook.)

We can call this kind of view, one on which collectives can be morally responsible due to their own agency-making characteristics, a non-reductive view of collective moral responsibility. Thus, some philosophers endorse the following:

Collective Moral Responsibility (CMR): Collectives can be irreducibly morally responsible.<sup>5</sup>

In this paper, I give a new argument against CMR. Although the argument is new, the idea underlying it is an old one: that collectives are not morally responsible because they do not have minds. My argument is an elaboration and defense of this basic idea. It does this in two ways; first, by providing a plausible alternative to theories which claim that collectives can have minds,<sup>6</sup> and second, by presenting two arguments that provide new support for the idea that having a mind is necessary for accountability. Readers who are interested in the former will want to focus on Sects. 2 and 3; readers who are only interested in the latter may want to skip to Sect. 4, although there is some reference in that section to earlier material.

Although these sections can be considered separately, they fit into an overall argument against CMR—an argument which avoids some of the controversial assumptions of previous arguments against CMR. For example, my argument does not rely on the claim that collectives are incapable of intentional states (c.f. Miller 2002). They possess them in a derivative way (see Sect. 4), but they are just as capable as us of having states that have the “aboutness” of intentionality.

It also does not assume that collectives lack some characteristics that have been thought important for moral responsibility. Let me explain. The rough idea behind CMR seems to be that collectives can be fully-fledged morally responsible agents, relevantly similar in all important respects to morally responsible individuals. In arguing that this is false, I will not be arguing that collectives share no morally important capacities with individuals; but I believe they lack some capacities required for full-fledged moral responsibility.

To clarify it is helpful to use David Shoemaker’s tripartite theory of moral responsibility.<sup>7</sup> Shoemaker distinguishes between three distinct notions which have been thought by different people to capture what we mean by “moral responsibility.” All three notions are tied, in a broadly Strawsonian way, to particular

---

Footnote 4 continued

nor those who had designed the informal notification system could be held responsible for the crash. One might think that the airline is responsible, but that this was not true of any individual (either within the airline at the time, or even those who designed the airline’s communications system). I owe this example to Phillips (1995, p 564 ff).

<sup>5</sup> Non-reductive moral responsibility has also been called non-distributive moral responsibility. Like most writing on this, I will assume that if collectives bear irreducible moral responsibility for their actions, it is because they are morally responsible agents who are related to those actions in the right way. For a recent (but, I think, unsuccessful) attempt to get collective non-reductive responsibility without collective agents, see Chant (2015).

<sup>6</sup> E.g. Theiner (2014).

<sup>7</sup> See Shoemaker (2015).

affective responses.<sup>8</sup> (Shoemaker argues that these three notions are distinct by arguing that some individuals are apt targets for some but not all of their associated responses.<sup>9</sup>) The first is attributability: one is attributable roughly because one is apt to be admired or disdained for one's character. While attributability has to do with character, answerability (the second notion) has to do with specific judgments: one is answerable roughly because one is apt to have others be disappointed or pleased at the quality of one's judgments. (Not only moral judgments—all sorts of judgments can be apt targets for an onlooker's disappointment, and so sufficient for answerability.)

The third notion is accountability, which will be of most interest here. One is accountable roughly because one is apt for others to have a particular sort of anger or gratitude towards one on the basis of one's quality of will. The sort of anger and gratitude in question is what Shoemaker calls *agential* anger and gratitude. For simplicity, we can focus on agential anger. This anger is an affective response to a (perceived) slight—a case where the object of your anger has disrespected you in some way—which is oriented towards making the offender realize what they have done to you.<sup>10</sup> Agential anger also sometimes—but not always—prompts us to seek revenge. (I am talking as though agential anger is always had on one's own behalf, which is not the case; one might feel similarly on behalf of someone else. We can set this detail aside for ease of expression.)

It is this agential anger—and related feelings—for which I will argue that collectives are not apt objects. I am thus denying that collectives are capable of accountability. This does not imply that they are incapable of attributability or answerability. In fact, I am happy to grant that collectives can have characters which make them apt targets for admiration or disdain, and can make good or poor moral judgments which make them apt for the pleasure and disappointment associated with answerability (this is plausibly a consequence of the sort of anti-aggregation arguments popularized by List and Pettit [e.g. in their (2011, chapter 2)]. Thus, I endorse collectives' capacity for two of Shoemaker's three senses of "moral responsibility."

Although accountability does not cover everything that one might mean by "moral responsibility," it is central to paradigmatic responsibility practices. Feelings which push us to confront someone who has wronged us are commonplace in daily life. And it is this sort of practice which is most clearly threatened by skepticism about moral responsibility. Dirk Pereboom, for example, who has prominently argued that we lack free will and should revise our practices in light of this, does not argue that we ought not think of some actions as expressions of an agent's deep self (attributability), or that some judgments are bad and others good

<sup>8</sup> In saying these notions of moral responsibility are connected to affective responses in a "broadly Strawsonian way," I mean that, for any of the three senses of moral responsibility, having one of them implies being an apt target for a certain set of affective responses. I do not mean to imply that being morally responsible in any of these senses metaphysically depends on our responses; that is a separate question (one I touch on when discussing premise 2).

<sup>9</sup> See Shoemaker (2015, chapters 1–3).

<sup>10</sup> For more discussion, see Shoemaker (2015, chapter 3), particularly pp 103–112.

(answerability). In taking our moral responsibility practices to be flawed, he is instead focused on our angry and sometimes vindictive response to wrongdoing.<sup>11</sup>

In saying that collectives are incapable of accountability, I am therefore arguing that they are incapable of being fit targets for attitudes that are paradigmatic parts of ordinary practices related to moral responsibility. But since previous writers on collective moral responsibility have not tended to distinguish these senses of moral responsibility (particularly answerability and accountability), it is not clear to me to what extent my position differs from theirs. It is possible that previous writers have simply not had accountability in mind when defending CMR, and instead have only intended endorsements of collective answerability. But I suspect that defenders of CMR are typically defending the notion that collectives are capable of anything one could reasonably mean by the term “moral responsibility,” which certainly includes accountability. If I am right about this, then defenders of CMR tended to assume that CMR implies the following:

Collective Accountability (CA): Collectives can be irreducibly accountable.

I will assume this implication as well, such that an argument against CA will also by implication target CMR. But even if I am wrong about this, establishing that collectives are incapable of accountability would be an interesting result in itself.

CA might strike some readers as crazy.<sup>12</sup> While I will argue that they are right to think that it is false, here is a reason to think that its falsity isn't obvious.<sup>13</sup> Large, complex, sophisticated organizations—collectives—play important roles in our lives. And these roles are morally charged: the way that collectives behave matters a great deal. Furthermore, social science and human experience suggest that collectives can transcend their members in morally important respects, if only because corporate cultures and structures can influence their members in profound ways—effects that are grounded in the behaviors of other individuals, but are not easily morally attributed to anyone in particular. This can lead to a sense that something of moral importance is lost when we blame individuals but do nothing to address the larger culture which influenced their behavior.<sup>14</sup> CA can be seen as a way of providing us with the something we were missing: a genuinely morally responsible agent who remains after we have dealt with all the individual wrongdoers. I suspect that the notions of collective attributability and answerability do at least as good a job of filling this moral void, and that CA is not required to justify

---

<sup>11</sup> Pereboom (2014, p 128).

<sup>12</sup> This is probably especially true for readers who are incompatibilists about moral responsibility—that is, who think that moral responsibility is incompatible with determinism. After all, it is unlikely that collectives possess the sort of capabilities that incompatibilists usually think is required for moral responsibility.

<sup>13</sup> In addition to the considerations I discuss here, some positions in philosophy of mind and action theory support the possibility of CMR; see Theiner et al. (2010).

<sup>14</sup> This is related to the problem of a “responsibility deficit” when individuals are punished for certain kinds of collective action. The basic idea is that punishing all the relevant individuals isn't enough; there is some responsibility “left over.” See Braham and van Hees (2011) for a recent discussion of the idea.

legal and social sanctions of collectives. But CA is, at least, an attempt to take the distinctively collective aspects of our lives morally seriously.

But since I am arguing that it is false, let's turn to the argument. Here it is:

1. If collectives are not capable of phenomenal consciousness, then collectives do not have minds.
2. If collectives do not have minds, then CA is false.
3. Collectives are not capable of phenomenal consciousness.
- ∴ 4. CA is false.

Since the argument is valid, I will spend the rest of this paper defending the premises. Since premise 3 is the least contentious and can be dealt with quickly, I will start with it and then move on to premises 1 and 2.

## 2 Premise 3

Premise 3 claims that collectives are not capable of phenomenal consciousness. In saying that something is phenomenal consciousness I simply mean that “there is something that it is like” to be that thing; that is, it has phenomenally conscious states, which are states that are experienced—states that involve what Chalmers calls “qualitative feels” (Chalmers 1996, p 4).<sup>15</sup> Premise 3 is closely related to a very non-controversial claim, that collectives do not in fact enjoy phenomenal consciousness. I know of no one who clearly denies this.<sup>16</sup> Premise 3 goes beyond this claim, since it claims not only that collectives are not phenomenally conscious, but also that they are not capable of phenomenal consciousness. As I am using the term “capable of phenomenal consciousness”, this means that collectives have also never been conscious, and cannot, in the normal course of things, become conscious. An unconscious human may lack phenomenal consciousness, but typically is capable of phenomenal consciousness; they once were conscious, and (perhaps with medical intervention) may reach that state again. I think this stronger claim should also be uncontroversial; clearly, corporations do not lack phenomenal consciousness just because they are asleep.<sup>17</sup>

<sup>15</sup> Note that, unlike Chalmers and many others, I am not committed to the view that “qualitative feels” lack representational or intentional content.

<sup>16</sup> Eric Schwitzgebel has suggested that, if materialism is true, then some collectives likely do enjoy phenomenal consciousness (Schwitzgebel 2015). But he isn't committed to the antecedent, and so doesn't endorse collective phenomenal consciousness. A referee has pointed to Schmid (2014) as a possible endorser of collective phenomenal consciousness. I am not sure I understand Schmid's position, but I do not think he is endorsing collective phenomenal consciousness. Even if he is, however, I take it that the considerations given in the next paragraph give us good reason not to accept such a view.

<sup>17</sup> One might think—some do, I believe—that collectives could possess phenomenal consciousness in fantastical scenarios, such as that envisioned by the Chinese Nation thought experiment, where the population of China acts in concert in such a way as to functionally mimic the activity of neurons. I take it, though, that exemplifying the sorts of states which would be required for phenomenal consciousness are vastly beyond the capabilities of any collective we are likely to create anytime soon. We can safely say, therefore, that actual and reasonably possible collectives are not capable of phenomenal consciousness. For simplicity, I leave this consideration aside in the text.

In addition to our strong intuitive sense that collectives are far from enjoying phenomenal consciousness, the falsity of premise 3 would cause problems for our normal ethical perspective on collectives. For if collectives were capable of phenomenal consciousness, then this would give us reason to think that they have moral interests of the sort held by animals who are capable of phenomenal consciousness. So, for example, we would have reason not to injure collectives since doing so might cause them pain. This makes premise 3 more than a merely theoretical matter. Insofar as we have reason to think our ordinary ethical perspective on collectives is approximately right, that gives us further reason to accept premise 3.

### 3 Premise 1

Premise 1 claims that phenomenal consciousness needs to be within a being's capacity if that being is to have a mind.

This has been defended by Uriah Kriegel and Terrance Horgan (Horgan and Kriegel 2008, pp 353–9). Their defense relies on a particular research program that Kriegel refers to as the “Phenomenal Intentionality Research Program,” or PIRP (see Kriegel 2013, p 1). Advocates of PIRP maintain that phenomenal consciousness has (often, at least) an intentional character, special features of which play an important role in a general understanding of intentionality. The sort of intentionality found within phenomenal consciousness is underived or basic intentionality; other intentional items (such as linguistic utterances) derive their intentionality in some way from this basic intentionality. (A typical example of the sort of special feature often claimed for phenomenal intentionality is non-derivatively determinate content; the idea is that only phenomenal states have determinate content in and of themselves (Kriegel 2013, p 10).) PIRP stands in contrast to the “Naturalist-Externalist Research Program” (NERP), which tries to understand intentionality by reference only to “natural” (typically causal/functional) features of the world, according no special place to phenomenal consciousness.

Kriegel and Horgan argue that PIRP, with its emphasis on the special role of phenomenal intentionality, fits well with the view that phenomenal intentionality is the “mark of the mental,” in this sense: only phenomenally intentional states are mental due only to their own nature, without any help from other states; and “other mental states count as mental only when, and insofar as, they bear the right relationship to phenomenally intentional states” (Horgan and Kriegel 2008, pp 253–4). The “right relationship” favored by Horgan and Kriegel makes reference to an ideal observer; if the observer would adopt an intentional stance towards the state, then the state is mental.<sup>18</sup> Although their official story is interpretivist in this sense, Horgan and Kriegel believe that an interpreter would in

---

<sup>18</sup> Their view here is similar to Dennett's; but it is worth noting that unlike Dennett's interpretivism, Horgan and Kriegel's interpretivism is not universal: some states (phenomenally conscious ones) are not dependent upon any observer's interpretation. This allows them to avoid some problems with Dennett's view; see Horgan and Kriegel (2008, p 355).

fact interpret as mental only states which bear the right sorts of causal relations to phenomenally conscious states.<sup>19</sup> What is really important about their view, for our purposes, is this claim that non-phenomenally conscious states are mental only because they are appropriately related to phenomenally conscious states.

To have a mind, on this view, is to be capable of phenomenally conscious states; non-phenomenally conscious mental states are always states *of* some mind with whose phenomenally conscious states they are appropriately related.<sup>20</sup> (One could imagine a mind as consisting of two parts: an inner core of phenomenally conscious states, with an outer layer of non-phenomenally conscious states that count as states of that mind just insofar as they are appropriately attached to the phenomenal core. Things that cannot have phenomenally conscious states lack a core for the outer layer to latch on to.) Furthermore, although all mental states might be intentional, not all intentional items are mental—it is perfectly possible to have intentional properties without being a part of any mind. Although these claims are not endorsed by all working within the PIRP program, for convenience I'll refer to their conjunction as the “PIRP view of minds.”

I believe that the PIRP view of minds is correct. This is partly because it provides an attractive way of delineating the mental from the non-mental. As Horgan and Kriegel argue, their view is superior in this respect to other theories of the mind, such as Brentano's idea that a state is mental just in case it is intentional, and Searle's idea that the mental consists of states which are either conscious or which could be conscious. Brentano's theory cannot account for non-mental intentional items like tokened sentences of natural language. Searle's view, although it rightly accords a central place to phenomenal consciousness, cannot account for the mentality of essentially unconscious states, that is, mental states for which it is psychologically impossible that they be conscious. It seems that many states associated with the visual system have this feature, such as states within the dorsal stream (Horgan and Kriegel 2008, pp 358–9).

The PIRP view of minds also gives the intuitively right results when confronted with cases. The sort of unconscious mental processes of normal humans that seem to clearly be mental (such as memories) are mental. Clearly non-mental things with

<sup>19</sup> In particular, the sort that support “cognitive, broadly inferential connections between the relevant states and phenomenally intentional states” (Horgan and Kriegel 2008, p 356). On this view, it is vague whether some states are mental or not, when those states are somewhat causally connected in the right way to phenomenally conscious states, but not strongly; see Horgan and Kriegel (2008, pp 355–6). Figuring out precisely what sorts of causal relationships would support such connections would be a major project of its own, but the details do not matter for our purposes; see Sect. 3.1, though, for discussion of an objection that this view—however the details are filled in—opens the door for collectives to have minds. I tentatively endorse their view on what sorts of relationships make non-phenomenally conscious things mental states. My endorsement carries a caveat: I am interpreting the causal integration at issue to include not only states which have actually had close causal relations, but also which would under the right circumstances; so a new brain state in an unconscious human could be causally related to phenomenally conscious states even before the human gained consciousness. It isn't clear to me whether Horgan and Kriegel use the term “causal integration” this liberally or not.

<sup>20</sup> This is not meant as a sufficient condition; perhaps it takes more than the capacity for phenomenally conscious states to be a mind. It is a necessary condition, one compatible with a number of different views about what minds are (e.g. properties of animals, immaterial substances, etc.).



intentional content (such as the written sentences which constitute this paper) are not, since intentional items are only mental states if they are phenomenally conscious or appropriately related to phenomenally conscious states. (Section 4 says a bit more about how this works on my view.)

Similarly, primitive creatures without phenomenal consciousness count as mindless, though possessing states whose functional roles mean they also have intentional content; so viruses and the like rightly come out as mindless. The view also does well when faced with more *recherché* examples. For example, Strawson's weather watchers [discussed in (Strawson 2010, p 251)], who have phenomenal experiences but neither the capabilities or dispositions for outward behavior are rightly counted as having minds. And functional zombies—those philosophical monsters that possess functional states duplicating our own but who lack phenomenal consciousness—very plausibly come out as mindless, *contra* purely functionalist views of the mental.<sup>21</sup> In fact, purely functionalist views of the mental—which are probably the best hope for the view that collectives can have minds, due to the functional sophistication of many collectives—probably also have significant difficulty dealing with the apparent facts that Strawson's weather watchers have minds and viruses do not.

Of course, space does not permit a full explication of the PIRP view of minds, let alone its place within PIRP or its overall merits relative to its rivals.<sup>22</sup> But I hope to have established that the PIRP view of minds has considerable plausibility.<sup>23</sup> Readers who withhold judgment about it may interpret the remainder of my argument in this section as a conditional: if the PIRP view of minds is true, then collectives do not have minds.

A reviewer has suggested to me that talking about minds in this context is unhelpful: talking about phenomenal consciousness is an advance in our theorizing that can replace the pre-theoretic term “mind.” For some projects, I think this kind of replacement is a good idea. But moral responsibility is also a pre-theoretic term, and when investigating concepts related to moral responsibility I think we need to be aware of possible connections between them and our concept of mind. Simply substituting talk of minds with the perhaps more precise talk of phenomenal consciousness isn't conducive to this.

One might object that, by introducing a theory about the meaning of “mind,” I can no longer appeal to connections involving our pre-theoretical view; but the PIRP of minds view I am defending is a theory about what our pre-theoretic view of

---

<sup>21</sup> Note that this does not necessarily imply that zombies lack intentionality, nor that zombies are genuinely possible or fully conceivable; and so everything I have said is consistent with physicalism, and indeed with the thesis that all phenomenal states are identical to states with certain functional properties. The idea is just that if (*per impossible?*) zombies existed, then they would not have minds.

<sup>22</sup> See, in addition to the papers already cited, (Horgan 2013).

<sup>23</sup> And to have established that this response (discussed in Theiner 2014, p 309) is inadequate: that views of mind that require phenomenal consciousness are “suspiciously anthropocentric” on the grounds that they “match those we take to indicate a *human* mind.” The PIRP view of minds does not arbitrarily pick some features of human minds and privilege them; its focus on the phenomenal is motivated by general views about intentionality, and it is capable of delivering plausible results about a wide range of non-human cases.

minds *actually is*, as well as what it ought to be. So, if true, the PIRP theory of minds should preserve connections between our pre-theoretical views of minds and of concepts related to moral responsibility.<sup>24</sup>

Here is another objection. The PIRP view of minds I am using was presumably developed by thinking about individual minds, and so one might think that we should not expect it to be applicable to collective minds. But I do not see any good reason to think that our pre-theoretical concept of mind is bifurcated between individual and collective cases in this way. The PIRP view is a strong candidate for an explication of our pre-theoretical view of minds, and gets a very wide variety of cases right; we should not arbitrarily limit its domain of applicability.

### 3.1 Does the extended mind cause problems for premise 1?

Here's an objection to what I've said thus far. I've said that there are reasons to be attracted to a PIRP view of minds, and that if we endorse this view, then we should also endorse premise 1. One might think that even if PIRP is correct, there is another flourishing research program that, while compatible with a PIRP view of minds, is incompatible with premise 1. I'm referring to the extended mind hypothesis. In this section, I will argue that this hypothesis, properly understood, poses no threat to premise 1.<sup>25</sup>

First, let's get a better handle on the objection. Consider Otto, whose reliance on his notebook has qualified him as a poster boy for the extended mind.<sup>26</sup> Otto's memory is poor, but by carefully recording facts in his notebook and carefully checking it, he can effectively remember all he needs to. According to the extended mind advocate, Otto's memories extend beyond his skull: they are partially realized by the words written in his notebook. Horgan and Kriegel have argued that the PIRP view of minds is consistent with this putative fact: since non-phenomenal states may be mental if they are causally connected in the right way to a phenomenal state, it may be that the connection between Otto's notebook states and his phenomenal states is close enough to count the notebook states as mental.<sup>27</sup> If so, it seems as though we have a mental state held by a larger system than Otto—the system of Otto + Otto's notebook—which suggests that this larger system has a mind.

<sup>24</sup> Of course, one could be a revisionist about moral responsibility, and think we should endorse CMR even though doing so doesn't accord with our pre-theoretic views about moral responsibility. In this case, any conceptual connections between our pre-theoretic notions might well be irrelevant. This would be a very interesting view, and perhaps one that deserves discussion; but as far as I can tell, it is not the view of the advocates of CMR, and I will set it aside.

<sup>25</sup> I'd like to thank Kate Ritchie for suggesting this objection to me.

<sup>26</sup> See Clark and Chalmers (1998) for an explication and defense of the extended mind that uses Otto as an example (on pp 12–13). This specific example has been criticized [see e.g. (Michaelian 2012)], but the details of these criticisms won't matter for our purposes, and Otto still functions very well as a rough guide to the extended mind hypothesis.

<sup>27</sup> As noted earlier (see fn. 12), Horgan and Kriegel suspect that there is “no deep fact of the matter as to whether [such states] are mental or not” (Horgan and Kriegel 2008, p 16). But they are open to the possibility that the mentality of Otto's notebook states are not indeterminate, and treating Otto's notebook states this way sharpens the objection, so we can set this complication aside.

Despite lacking the capacity for its own distinctive phenomenal consciousness, it is appropriately related to Otto's, and so can fulfill the PIRP view of mind's requirement that all mental states be appropriately related to phenomenal states. And if this is the case for Otto + his notebook, mightn't it be the case for collectives as well?

We can reconstruct this objection in the following way:

5. If the PIRP view of minds is true, then, possibly, the extended mind hypothesis is true.
6. If the extended mind hypothesis is true, then the phenomenal intentionality of an individual can partially ground the fact that a larger system S has a mind.
- ∴ 7. If the PIRP view of minds is true, then, possibly, the phenomenal intentionality of an individual can ground the fact that a larger system S has a mind. (5, 6)
8. If (7) is true, then the PIRP view of minds does not give us any very good reason to endorse premise 1.
- ∴ 9. The PIRP view of minds does not give us any very good reason to endorse premise 1. (7, 8)<sup>28</sup>

This argument is unsound. According to premise 6, the Extended Mind hypothesis implies that some system S, composed of some individual's internal states (e.g. their brain) and some external state (such as Otto's notebook), but *other than the individual* (e.g. other than Otto), has a mind.

I think this is a misreading of the extended mind hypothesis. It is best seen (and, I think, typically seen) as about what instantiates the mental states of individuals like Otto. Does Otto's mental life end at the skin, or do external states of affairs partly instantiate Otto's mental states? The extended mind hypothesis claims that the latter is the case: Otto's mind extends beyond his skin. It does not claim, for the case of Otto, that Otto's mind is restricted to his skin but that Otto + his notebook possesses its own mind. Why should it? It is more plausible to think that we have a single, extended mind—Otto's—then to think that there is some second mind out there in addition to Otto's but existing by virtue of his phenomenal consciousness.

Of course, Otto's case is a simple one; collectives are far more complex and sophisticated than Otto and his notebook. But I cannot see how that should make any difference here, at least assuming that the PIRP view of minds is correct. Collectives may have properties which are appropriately related to the phenomenally conscious states of individuals, but that only implies that these are the mental states of the individuals with which they are related, not that they are mental states of some new mind.<sup>29</sup> The extended mind hypothesis does not give us any reason to think such additional minds exist. (Of course, one might think that the complexity of

<sup>28</sup> Some of the language of this reconstruction is due to Kate Ritchie.

<sup>29</sup> I should point at that on the view I am defending, it is quite possible for something to be the instantiation of two different mental states held by two different minds. For example, we could imagine that Otto's notebook is used in just the same way by Ottolina. In that case, both Otto's and Ottolina's memories would be partially instantiated within the notebook. The same letters on the page partially instantiate different mental states with the same contents. If the extended mind hypothesis is correct, then

collectives gives us reason to think they have minds for some other reason, such as their functional sophistication; but that would involve a rejection of the PIRP view of minds, and not be a version of the objection that the conjunction of the PIRP view of minds and the extended mind hypothesis causes a problem for premise 1 of my argument.)

So premise 6 is false; there is no connection between the extended mind hypothesis and collective mentality. The extended mind hypothesis thus poses no problem for my argument.

In this section I have set out a view about what it takes to have a mind—the PIRP view of minds—which has not received much attention in discussions of CMR. I take this section, in conjunction with the section on premise 3, to give us good reason to think that collectives do not have minds. At least, we should believe this if we accept the PIRP view, a view which has a good deal to recommend it.

## 4 Premise 2

Premise 2 claims that collectives cannot be irreducibly accountable if they do not have minds. Calling something “mindless” seems to put it out of the running for accountability. There is thus, I take it, some initial intuitive support for premise 2.

One might object to premise 2, however. Suppose that I am correct that collectives lack minds. They do, nonetheless, possess states which functionally imitate mental states, and which may in fact imitate them exactly with regard to their functional roles.<sup>30</sup> Why not think that these pseudo-mental states, which play the same functional roles as the genuine articles, also play the same roles with respect to securing irreducible accountability?<sup>31</sup> Presumably, they would do this because their similarity of functional role secures similar intentional content. Call this the *Same Intentional Content Objection* (or—with my apologies—SICO).

In light of this objection, the basic intuition described above needs to be defended; ideally, such a defense would explain what it is about collective’s lack of minds that make them incapable of accountability. If the PIRP view of minds is correct, then it is plausible that such an explanation will make reference to collective’s lack of capability for phenomenal consciousness. I will offer two such defenses here.

The first starts from the fact that accountability requires being apt for attitudes that have a communicative goal: that of making the offender aware of how they have slighted us.<sup>32</sup> We do not want a mere apology coupled with a change in future

---

Footnote 29 continued

something like this is probably often the case with collectives; a shared office calendar would be an obvious simple example.

<sup>30</sup> For discussion of this, see e.g. (Copp 2006) and Hess (2010, 2013).

<sup>31</sup> I think this is the rough idea behind many contemporary arguments for CMR; see, for example, the papers from the previous footnote.

<sup>32</sup> Here, as elsewhere, I am ignoring the positive analogue to agential anger, agential gratitude. This should not matter for the argument.

behavior. We want to make them experientially understand (to some degree) just what it was like to be us as the slighted party: “[we want the offender to be] made to appreciate, as fully as possible, what it is that he did to [us], and he cannot be gotten to appreciate that without having been gotten to *feel what it was like* when he did it” (Shoemaker 2015, p 110; emphasis in original). That is, the kind of appreciation we are after requires that the offender feel a certain sort of way. (Consider the anger you feel when you have been treated in an insulting or hurtful manner: you probably want the person who treated you this way to experience a kind of empathic response that makes them experience something of what your distress was like, and your expression of anger towards them is supposed to provoke this response.)

As Shoemaker notes, the fact that our agential anger is directed towards making our target *experientially* understand the way they have hurt us implies that agents whose capacity for sharing our feelings is inhibited have only a mitigated capacity for accountability (Shoemaker 2015, 171). They are less than fully apt as targets for the reactive attitudes accountability is about. Shoemaker was addressing the subject of humans with significantly impaired empathic capabilities, but who nonetheless are phenomenally conscious. Collectives, though, are incapable of phenomenal consciousness at all. They therefore cannot have any sort of experiential understanding. This means that they are not apt targets for our agential anger, and hence that they are not accountable.

When considering this argument, it is important to bear in mind that its conclusion is not that there are no responses to collective wrong-doing that make sense. If a corporation does something wrong, not only are a variety of sanction-oriented actions possibly appropriate, but so also are intense moral disappointment and disapproval, perhaps even moral outrage that such a wrong-doing occurred. [Some of these responses, such as disapproval, are explicitly licensed by Shoemaker’s theory under the notion of answerability—bad actions are worth criticizing, after all (Shoemaker 2015, chapter 2)]. And, of course, the individuals who make up the group may also be apt targets for all sorts of responses. So there are still a number of tools in our moral toolbox for dealing with cases of collective wrong-doing; but by the same token, the existence of these tools does not give us reason to think that premise 2 is false.

It is instructive at this point to look at a recent attempt to argue that phenomenal consciousness is not required for moral responsibility. Björnsson et al. (2016)—who are also working within a broadly Strawsonian tradition focused on our attitudes and practices—have argued that collectives can have capacities that are the moral equivalents of reactive attitudes. If so, perhaps corporations can avoid the need for phenomenal consciousness I just claimed.

In response, I argue that Björnsson and Hess have shown only that collectives can be capable of answerability, not accountability.<sup>33</sup> They argue that one might have several reasons for thinking that phenomenal consciousness is necessary for moral

---

<sup>33</sup> Björnsson and Hess do not distinguish between answerability and accountability, talking instead about moral responsibility. As I noted in the introduction, I think that they intend for their account to cover everything one could plausibly think important to moral responsibility, including what I call accountability. But even if they do not, their work could be adapted to provide the objection considered here.

responsibility; chiefly, that guilt involving phenomenal pain might be necessary to motivate non-strategic apologies and other changes in behavior, or to make the agent aware that they have acted badly. They reasonably argue that phenomenal pain is unnecessary for these motivational and epistemic functions; collectives can have functional analogues of phenomenal pain—a sort of “internal disorder”—that play these roles. This shows that something does not need phenomenal consciousness to be able to fulfill these functions (see Björnsson et al. 2016, pp 20–25).

However, the functions that Björnsson and Hess consider are not sufficient for accountability; a collective could have the capacity for them merely by being answerable (i.e. capable of good or poor decisions). This can be seen by the fact that these functions can be involved in cases that do not involve moral failures at all, let alone agential anger. (Recall that one can be answerable for non-moral judgments as well as moral ones.) Think about a stereotypical gym teacher overseeing some organized sports, publicly shaming students when he judges that the embarrassment will help them focus on the ways that their performance does not measure up, and give them reason to apologize to their teammates and perform better in the future. [Of course, that does not mean that causing pain in such a situation for those reasons would be morally permissible, or pedagogically effective. Even if it is neither, this does not mean that the student is not an apt target for such responses in the relevant respect; for discussion of similar issues see (Shoemaker 2015, pp 106–7.)]

The fact that all these actions make sense in a non-moral context shows they miss something distinctive about agential anger; in particular, as I suggested above, that agential anger motivates us to cause pain (partly) so that the pain itself be part of their appreciation of what it was like for us to suffer due to their actions. It is plausible that such an appreciation would need a phenomenal component; in any case, Björnsson and Hess’ account of the motivational and epistemic power of collective analogues of pain does not give us reason to think otherwise. In sum, then, Björnsson and Hess’ case might suffice to show that collectives can be answerable (which, given their lack of minds, might be itself a surprising and interesting result), but does not show that collectives can be accountable, and so does not cause a problem for my argument.

So much for my first defense against the idea that the Same Intentional Content Objection gives us good reason to abandon the intuitive “mind” requirement; in my second defense I will argue that accepting SICO leads to an implausible consequence. Before we get there, however, we need to return to PIRP—the phenomenal intentionality research program—and the sort of views of intentionality PIRP has fostered. As noted above (in my discussion of premise 1), advocates of PIRP think that phenomenally conscious states (at least sometimes) possess a sort of basic, original, or underived intentional content, and that all other intentional states are intentional in some non-basic or derived way. Now, although I’ve said a few words earlier in defense of what I called the PIRP view of minds, I have not defended this related view about intentionality. Nor will I do so here, since it would take as very far afield indeed.<sup>34</sup> But readers suspicious of this approach to

<sup>34</sup> Readers interested in arguments for this view of intentionality might be interested in Searle (1991, 1992), Horgan and Tienson (2002), and Kriegel (2003).

intentionality will, I suspect, also reject the PIRP view of minds I sketched earlier; and so I hope that those who remain think this is at least a plausible view about intentionality. I will assume it from here on out.

Now, PIRP theorists owe us an account of how it is that things with derivative intentionality get that intentionality.<sup>35</sup> Different theorists give different answers.

Some (following Searle) think that it is by being potentially phenomenally conscious, others (such as Loar and Horgan) that it is by having inferential connections to phenomenally conscious states; and finally Kriegel has developed an interpretivist view on which states with derivative intentionality have them by virtue of the fact that an idealized observer would interpret them appropriately.<sup>36</sup> The chief difficulty with the first two views is that they get the extensionality of intentionality wrong. Neither can account for intentionality that is not part of a mind, such as that of written words; words on a page are not themselves potentially phenomenally conscious, nor are they inferentially related to phenomenally conscious states. Furthermore, as Kriegel discusses, they have problems accounting for the intentionality of some cognitive states such as those in the dorsal stream of the visual cortex (see Kriegel 2011, pp 192–198). And, in any case, none of these other theories can account for the apparent intentionality of collectives, since collective's intentional states are not potentially phenomenally conscious, nor always inferentially related to phenomenally conscious states. So interpretivism of Kriegel's type is the best account of which I am aware.<sup>37</sup>

On this view, non-phenomenally intentional objects have their intentionality by virtue of the fact that an idealized observer would find it pragmatically useful to count such objects as intentional, using Dennett's "intentional stance." (And their specific intentional content would be had by virtue of the fact that the observer would attribute that specific content to them.)<sup>38</sup> I think this theory is independently plausible as an account of the nature of collective intentionality, apart from any more general theoretical merits. It accounts for what seems to me to be a general tendency (among philosophers and ordinary folk) to resist ascribing intentionality to a collective's states until it is practically useful to do so, and to regard this practical

<sup>35</sup> This is a different issue from that of what makes something unconscious a mental state. Something could be intentional without being a mental state. So although some of the moves made to account for derivative intentionality are similar to the sorts of accounts discussed earlier about what makes unconscious things mental states, they have a different subject matter that could in principle require different stories.

<sup>36</sup> See Kriegel (2011, pp 189–206) for discussion. Kriegel also discusses eliminativism, which I will set aside as getting the extensionality of intentionality even more wrong than the first two views.

<sup>37</sup> One might think that we should abandon Kriegel's interpretivism for an account on which derived intentionality is held by virtue of whatever his ideal interpreter would use as a basis for their judgments. A key reason not to is explanatory simplicity: the basis on which the interpreter judges is far more complex (in Kriegel's term, "heterogenous") than the judgments themselves, which arguably gives us reason to think that it is the judgment that is doing the explanatory work. See Kriegel (2011, pp 216–217) for more discussion. James Otis has suggested to me that actual observers might be able to do some or all of the work that Kriegel attributes to a possible, idealized observer. This is an interesting suggestion, but I do not think that whether it is correct or not will bear on my purposes in this paper.

<sup>38</sup> Tollefsen (2002) has an interpretivist theory along these lines, although it is not developed within a PIRP framework.

usefulness as a good, even decisive, reason to think that the collective's states are in fact intentional. So, according to a plausible development of PIRP, the intentional properties of creatures without phenomenal consciousness—and collectives in particular—are response dependent properties. Assuming, then, that this is correct, here is how it tells against SICO.<sup>39</sup>

I think it is implausible to hold that the intentional properties required for accountability are response dependent properties of this sort. Something is not accountable simply because it is pragmatically useful to treat it as accountable. By way of analogy, consider the objection that many have to what has been called “metaphysical Strawsonianism,” the view that accountability is grounded in our practices of holding accountable.<sup>40</sup> It is natural to think that metaphysical Strawsonianism is incorrect, since our accountability practices presuppose facts about being accountable.<sup>41</sup> In both cases we are being asked to think of something fundamental to moral responsibility as being grounded in something about the thoughts and attitudes of others, rather than of the putatively accountable subject themselves. This gets things the wrong way around. The judgments of an idealized observer about accountability are not prior to the phenomena themselves.<sup>42</sup>

<sup>39</sup> One might think that contemporary work on the intentional states of collectives in particular, such as List and Pettit (2011) and Hess (2014), provides an alternative to the view that these states are response dependent properties. Perhaps this is true, but I am skeptical. Specific accounts of how particular entities (such as collectives) have intentional states need to fit into overall general views about the nature of intentionality (it is implausible to think that each domain of derived intentionality has its own account with no overall theory that provides a general account). And general views of the sources of intentionality that are offered or supposed by List, Pettit, and Hess make no mention of phenomenal intentionality, and fit far better in a NERP framework than a PIRP one. (Recall that PIRP views of intentionality ground non-phenomenal intentionality in phenomenal intentionality, while NERP views do not.) List and Pettit's functionalist views of intentionality, for example, do not leave any role to be played by phenomenal consciousness. Hess, as best I can tell, assumes a broadly functionalist view as well. So, on PIRP, these theories are incomplete, since they lack an account of how collective intentionality is grounded in phenomenal intentionality; and the most plausible ways I know of completing them seem to me (for reasons described in the text) to make them response-dependent properties. Thanks to an anonymous reviewer for suggesting that I address these theories here.

<sup>40</sup> For a helpful recent discussion of the difficulties with metaphysical Strawsonianism, see McKenna (2012, pp 31–55). Typical discussions of metaphysical Strawsonianism (including McKenna's) put the view in terms of moral responsibility, rather than accountability, but I think that the point is most convincing when it is accountability in particular we are considering, due to accountability's inherently confrontational nature (it would be strange if the intentionality of behavior warranting such confrontation were most fundamentally grounded in our practice of confrontation rather than in the agent itself).

<sup>41</sup> One might be concerned that this is in tension with my earlier defense, which relied on Shoemaker's theory; Shoemaker himself seems to be a metaphysical Strawsonian. But that is not essential to his theory, and I do not believe anything I have claimed is inconsistent with rejecting that view. For Shoemaker's discussion of metaphysical Strawsonianism as detachable from the rest of his views, see Shoemaker (2015, pp 20–21).

<sup>42</sup> One area of recent concern for moral responsibility theory has been responsibility for actions stemming from our unconscious mental states. (Actions stemming from implicit bias are one example.) It is controversial whether and how these states play a role in our moral responsibility, and so one natural question to ask about the argument I just gave is whether it is compatible with thinking that individuals are responsible for such actions. One might think that what I have said implies that we cannot be accountable for such behavior, because it implies that our unconscious mental states are only intentional in a response dependent way. I am not sure, however. I do not have the space to give this a full discussion,



The upshot of this section is that we have reason to think that actual minds—not simply their functional mimics—are necessary for accountability, since only actual minds have the phenomenal consciousness and underived intentionality required for accountability. At least, those who think that our accountability practices require the capacity for an empathic response, or who agree that SICO leads to an implausible response dependence for the capacities required for accountability, should think so.

Shoemaker's separation of accountability from answerability is important for the first argument in this section, and theories of intentionality are important for the second. This suggests that getting clearer on the issue of collective moral responsibility may require more attention to issues in the philosophy of mind, and to questions about the extent to which "moral responsibility" is even a unified phenomenon. Although there has been some work on the former recently (e.g. Hess 2010), neither topic has taken center stage in most contemporary work on collective moral responsibility. If nothing else, I hope that even those who are not convinced premise 2 is true will think that progress in this area relies partly on more attention to these subjects.

## 5 Conclusion

One natural set of claims opponents of CMR are likely to want to make is that moral responsibility requires having a mind, and that collectives do not have minds. These claims have come under attack, but there is reason to think that they have more life in them than has sometimes been thought. In particular, PIRP, a flourishing program in the philosophy of mind, offers reason to think that collectives lack minds, while attention to the variety of practices that play important part in our moral lives allows us to acknowledge collective's important similarities to us while recognizing that they are not capable of accountability.

**Acknowledgements** I would like to thank Brett Sherman, Bill FitzPatrick, Kate Ritchie, Dustin Crummett, Bill Rowley, Kelly Annesley, James Otis, Hayley Clatterbuck, an anonymous reviewer for *Philosophical Studies*, and an audience at the Pacific APA for helpful feedback and discussion.

---

Footnote 42 continued

but I think that our unconscious mental states may play roles in our moral responsibility for other reasons than their intentionality. It may be, for example, that the moral relevance of implicit racism has to do with its effects on our behavior, combined with our negligence or inability in our phenomenally conscious control over such effects, and not from any intentionality the implicitly racist states may possess. One might also hold a mixed view about derivative intentionality, on which some states—those suitably connected to phenomenal consciousness—are intentional by virtue of those causal connections, while others—including those had by collectives—are intentional by virtue of what an ideal observer would say about them. Of course, motivating such a disjunctive view might be a challenge, but the possibility of such a view makes me hesitant to draw conclusions about ordinary individual unconscious states from my argument in the text above.

## References

- Björnsson, G & Hess, K. (2016). Corporate crocodile tears? On the reactive attitudes of corporate agents. *Philosophy and Phenomenological Research*. doi:[10.1111/phpr.12260](https://doi.org/10.1111/phpr.12260)
- Braham, M., & van Hees, Martin. (2011). Responsibility voids. *The Philosophical Quarterly*, 61 (1), 7–15.
- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. New York: Oxford University Press.
- Chant, S. R. (2015). Collective responsibility in a hollywood standoff. *Thought*, 4, 83–92.
- Clark, A., & Chalmers, David. (1998). The Extended Mind. *Analysis*, 58(1), 7–19.
- Copp, D. (2006). On the agency of certain collective entities: an argument from ‘normative autonomy’. *Midwest Studies in Philosophy*, 30(1), 194–221.
- Hess, K. M. (2010). The modern corporation as a moral agent: the capacity for ‘thought’ and ‘first-person perspective’. *Southwest Philosophy Review*, 26(1), 61–69.
- Hess, K. M. (2013). The free will of corporations (and other collectives). *Philosophical Studies*, 168(1), 241–260. doi:[10.1007/s11098-013-0128-4](https://doi.org/10.1007/s11098-013-0128-4).
- Hess, K. (2014). Because they can: The basis for the moral obligations of (certain) collectives. *Midwest Studies in Philosophy*, 38, 203–221.
- Horgan, T. (2013). Original intentionality is phenomenal intentionality. *The Monist*, 96(2), 232–251.
- Horgan, T., & Kriegel, U. (2008). Phenomenal intentionality meets the extended mind. *The Monist*, 91(2), 347–373.
- Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. J. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings*. New York: Oxford University Press.
- Kriegel, U. (2003). Is intentionality dependent upon consciousness? *Philosophical Studies*, 116, 271–307.
- Kriegel, U. (2011). *The sources of intentionality*. New York: Oxford University Press.
- Kriegel, U. (2013). The phenomenal intentionality research program. In U. Kriegel (Ed.), *Phenomenal intentionality* (pp. 1–26). Oxford: Oxford University Press.
- List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford: Oxford University Press.
- McKenna, M. (2012). *Conversation and responsibility*. New York: Oxford University Press.
- Michaelian, K. (2012). Is external memory memory? Biological memory and extended mind. *Consciousness and Cognition* 21(3), 1154–65. Elsevier Inc. doi:[10.1016/j.concog.2012.04.008](https://doi.org/10.1016/j.concog.2012.04.008).
- Miller, S. (2002). Against Collective Agency. In Georg Meggle (Ed.), *Social facts and collective intentionality* (pp. 273–298). New York: German Library of Sciences.
- Pereboom, D. (2014). *Free will, agency, and meaning in life*. Oxford: Oxford University Press.
- Phillips, J. (1995). Corporate moral responsibility: When it might matter. *Business Ethics Quarterly*, 5(3), 555–576.
- Schmid, H. B. (2014). Plural self-awareness. *Phenomenology and the Cognitive Sciences*, 13, 7–24.
- Schwitzgebel, E. (2015). If materialism is true, the united states is probably conscious. *Philosophical Studies*, 172(7), 1697–1721.
- Searle, J. R. (1991). Consciousness, unconsciousness, and intentionality. *Philosophical Issues*, 1, 45–66.
- Searle, J. R. (1992). *The rediscovery of mind*. Cambridge, MA: M.I.T. Press.
- Shoemaker, D. (2015). *Responsibility from the margins*. Oxford: Oxford University Press.
- Strawson, G. (2010). *Mental reality* (2nd ed.). Cambridge, MA: MIT Press.
- Theiner, G. (2014). A beginner’s guide to group minds. In Mark Sprevak & Jesper Kallestrup (Eds.), *New waves in philosophy of mind* (pp. 301–322). New York: Palgrave MacMillan.
- Theiner, G., & O’Connor, T. (2010). The emergence of group cognition. In A. Corradini & T. O’Connor (Eds.), *Emergence in science and philosophy* (pp. 78–117). Abingdon: Routledge.
- Tollefsen, D. (2002). Collective intentionality and the social sciences. *Philosophy and the Social Sciences*, 32, 25–50.