

## *Individual and Collective Responsibility*

Andrew C. Khoury  
andrew.khoury@asu.edu  
Arizona State University  
Tempe, AZ  
USA

**Abstract:** Building on Peter French's important work, this chapter draws three distinctions that arise in the context of attributions of moral responsibility, understood as the extent to which an agent is blameworthy or praiseworthy. First, the subject of an attribution of responsibility may be an individual agent or a collective agent. Second, the object of the responsibility attribution may be an individual action (or consequence) or a collective action (or consequence). The third distinction concerns the temporal dimension of the responsibility attribution. Sometimes responsibility for an action is attributed to an agent at the time of the action. At other times responsibility for an action is attributed to an agent sometime after the action has taken place. Taken together, these three binary distinctions yield eight types of responsibility attribution. It is argued that a collective agent's responsibility for a past collective act is properly understood on the same theoretical model as is an individual's responsibility for a past individual act. While most assume that responsibility over time is a straightforward matter of identity over time, it is argued that instead this is a matter of psychological or attitudinal connectedness. The possibility is considered that this relation also grounds attributions involving an asymmetry between subject and object, such as individual responsibility for past collective action, but a skeptical worry is raised that such attributions entail an unpalatable form of moral luck and should therefore be rejected.

**Keywords:** Blameworthiness; Collective responsibility; Diachronic responsibility; Moral luck; Personal identity.

## **Introduction**

Like most philosophically interesting notions, attributions of responsibility and blame are far from straightforward. This is due, in part, to the fact that there are a number of different senses of such notions to which one might appeal. For example, Peter French argues that the following claim is not necessarily inconsistent:

Charles Manson is to blame for the Tate/LaBianca murders, but he cannot be blamed, even though we do blame him (1976, 43).

Roughly, the first occurrence of “blame” functions to make a purely causal claim, the last refers to our actual emotional responses, and the second concerns whether such responses are warranted. It is worth carefully distinguishing amongst these senses. But even when we are focused on a particular sense of responsibility or blame, such as the extent to which blaming attitudes are warranted, there are still a variety of notions that we may have in mind when making a responsibility attribution. My task in this chapter is to draw three distinctions that are applicable when discussing moral responsibility, in terms of blameworthiness and praiseworthiness, in the context of both individuals and collectives. Peter French’s work on responsibility is deserving of its influence and the claims developed and defended here, as will become evident, owe much to this important work.

The three binary distinctions drawn give rise to eight types of responsibility attribution. I will discuss each type in turn, but given the wide scope of the topic what I will say about each must be fairly brief. My chief purpose, then, is more that of the surveyor than that of the architect, though I will attempt to do some theory building.

The sense of responsibility with which I will be concerned is that of moral responsibility. By moral responsibility, I mean the extent to which one is blameworthy or praiseworthy. As noted above, this sense of responsibility should be distinguished from mere causal responsibility (e.g. “the rain is responsible for ruining your shoes”) as well as the question of whether a given agent is actually held responsible by anyone (e.g. “the electorate held him responsible for the scandal”). It should also be distinguished from legal responsibility (e.g. “she was found responsible for misconduct”) and responsibility in the sense of duty or obligation (e.g. “it is his responsibility to watch the children”). The sense of responsibility that I will be discussing is the sense in which being responsible is a matter of being an apt target of the reactive attitudes such as guilt, resentment, and indignation on the one hand and (certain forms of) pride, gratitude, and approbation on the other.<sup>1</sup> One important feature of this sense of responsibility is that it admits of degrees. One can, for example, be more or less blameworthy or more or less praiseworthy.

### **Three Distinctions**

An attribution of moral responsibility, in the intended sense, typically has the following form (or readily submits to such a form):

Subject S is morally responsible (i.e. blameworthy or praiseworthy) to degree  $d$  for object O.

---

<sup>1</sup> See Strawson 1962.

Such attributions involve a three-place predicate relating a subject and object to some degree or other.<sup>2</sup> Much attention has been directed at such attributions in the individual context. These are contexts in which both the subject is an individual agent and the object is an individual action (or the consequence of some individual action).

In the latter half of the 20<sup>th</sup> century, and owing, in part, to the great harms that occurred during WWII and the Vietnam War, some theorists turned their attention to responsibility in a collective context. For example, French argued that there are certain kinds of actions that can only be committed by collectives, not individuals:

There is, of course, a class of predicates that just cannot be true of individuals, that can only be true of collectives. Examples are abundant, and surely include “disbanded” (most uses of), “lost the football game”, “elected a president”, and “passed an amendment.” (1998, 37)

What French was suggesting is that in addition to individual actions, there is a large and significant class of actions (and corresponding consequences) that can only be brought about collectively. French argued further that not only are there collective actions, but there are also collective agents that are not straightforwardly reducible to individual agents.<sup>3</sup> For now, simply note the conceptual possibility that in a responsibility attribution of the form “subject S is responsible to degree  $d$  for object O,” the subject may be an individual or a collective agent and the object may be an individual action (or

---

<sup>2</sup> See French 1976, 443-444. He argues that blame is a two-place predicate (e.g. “We blame X for Y”). Insofar as one can be more or less blameworthy for a given object, responsibility in the sense of blameworthiness or praiseworthiness is (at least) a three-place predicate.

<sup>3</sup> See French 1984.

consequence) or a collective action (or consequence).<sup>4</sup> I am not now making a claim about the truth or falsity of such attributions, I mean only to set them out for subsequent discussion.

In addition to the distinctions concerning subject and object there is another distinction that concerns the temporal dimension of a responsibility attribution. Sometimes we attribute moral responsibility for an action to an agent at the time of the action, what I call *synchronic responsibility*. At other times we attribute responsibility for an action to an agent at some time after the action has occurred, what I call *diachronic responsibility*.<sup>5</sup> More precisely, synchronic responsibility concerns the extent to which an agent at time  $t_1$  is responsible for an action that occurs at  $t_1$ . Diachronic responsibility concerns the extent to which an agent at some later time  $t_2$  is responsible for an action that occurred at  $t_1$ .<sup>6</sup> Synchronic responsibility involves the responsibility of an agent at the time of action, while diachronic responsibility involves the responsibility of an agent at some time after the action occurs.

We have, then, three binary distinctions concerning attributions of moral responsibility, understood as the extent to which an agent is blameworthy or praiseworthy:

---

<sup>4</sup> For ease of expression I'll speak of the object as being an action rather than an action or consequence, but the disjunction should be taken as implied.

<sup>5</sup> See Khoury 2013; also see Matheson 2014 who uses the language of synchronic and diachronic ownership.

<sup>6</sup> The point here is that moral responsibility needs to be indexed to time. The manner in which I've placed the temporal index here naturally accords with a four-dimensionalist ontology; the subject of a responsibility attribution is an agent *at a time* rather than an agent *simpliciter*. But this is not the only option, and nothing in what follows requires the adoption of four-dimensionalism. For example, one might instead hold that moral responsibility is a four-place rather than three-place predicate relating subject, object, degree, and time. Or one might hold that instantiation itself is a temporal relation or appeal to a temporal sentential operator. Which account one prefers will depend on one's favored solution to the problem of temporary intrinsics (see Lewis 1986, 203-205).

(a) **Subject:** Individual agent or collective agent.

(b) **Object:** Individual action or collective action.

(c) **Temporality:** Synchronic or diachronic.

Taken together these three distinctions yield eight different types of moral responsibility attribution. These are represented in Table 1. “S” denotes the subject, “O” denotes the object, and the subscript denotes individual (i) or collective (c). I discuss these types of responsibility attribution below.<sup>7</sup>

**Table 1** Types of moral responsibility attribution

Type	Subject	Object	Temporality	Formal Structure
1	individual	individual	synchronic	$S_i$ at $t_1$ is morally responsible to degree $d$ for $O_i$ at $t_1$ .
2	individual	individual	diachronic	$S_i$ at $t_2$ is morally responsible to degree $d$ for $O_i$ at $t_1$ .
3	collective	collective	synchronic	$S_c$ at $t_1$ is morally responsible to degree $d$ for $O_c$ at $t_1$ .
4	collective	collective	diachronic	$S_c$ at $t_2$ is morally responsible to degree $d$ for $O_c$ at $t_1$ .
5	individual	collective	synchronic	$S_i$ at $t_1$ is morally responsible to degree $d$ for $O_c$ at $t_1$ .
6	individual	collective	diachronic	$S_i$ at $t_2$ is morally responsible to degree $d$ for $O_c$ at $t_1$ .
7	collective	individual	synchronic	$S_c$ at $t_1$ is morally responsible to degree $d$ for $O_i$ at $t_1$ .
8	collective	individual	diachronic	$S_c$ at $t_2$ is morally responsible to degree $d$ for $O_i$ at $t_1$ .

<sup>7</sup> Note that on this conception the distinction between individual responsibility and collective responsibility is wholly accounted for in terms of the relata of the moral responsibility relation. Alternatively, one might attempt to account for the distinction, not in terms of distinct relata, but in terms of distinct relations as May 1987, ch. 4 appears to hold.

## **Type 1**

Attributions of Type 1, individual synchronic responsibility for individual action, are familiar (though not necessarily under that description). In such attributions a claim is made about the extent to which an individual is responsible for an individual action at the time the act occurs. For example, suppose that Red has committed murder at time  $t_1$ . An attribution of Type 1 makes a claim about the extent to which Red at  $t_1$  is blameworthy for his murderous action that occurs at  $t_1$ .

The debate between compatibilists and incompatibilists about causal determinism and moral responsibility is best understood, I think, as a debate most fundamentally about the conditions of Type 1 attributions of moral responsibility. Incompatibilists argue that the truth of Type 1 attributions is incompatible with determinism while compatibilists deny this. There are many accounts of Type 1 in the literature and much of the debate has concerned the kind of control that is required for the truth of such attributions. Some argue that this type of responsibility requires agent-causal freedom, others argue that it is simply a matter of the proper alignment amongst one's desires, while others argue that it is a matter of being properly responsive to reasons.<sup>8</sup> For now let us note that while much of the debate has focused on the control condition, all plausible accounts will admit that there are conditions relating to the agent's beliefs and desires at the time of action.<sup>9</sup> This is simply the rather trivial claim that whether and the extent to which an (individual) agent is (synchronically) blameworthy or praiseworthy for an action is partly determined by the beliefs and desires from which she acted. For example, if Red met the proper

---

<sup>8</sup> See, for example, O'Connor 2000, Frankfurt 1971, and Fischer and Ravizza 1998.

<sup>9</sup> For example, it is widely acknowledged that there is an epistemic condition on responsibility.

control condition with respect to his killing of the victim, then his Type 1 blameworthiness will be largely a matter of his beliefs about what he was doing and the quality and content of the values and desires that moved him to act.

## **Type 2**

Attributions of Type 2, individual diachronic responsibility for individual action, are common in our lives though less commonly discussed explicitly in the literature. An attribution of Type 2 involves making a claim about the extent to which an individual at a later time is responsible for some action that occurred at an earlier time.<sup>10</sup> For example, suppose that Red has committed murder at time  $t_1$ . An attribution of Type 2 makes a claim about the extent to which Red at some later time  $t_2$  is blameworthy for his murderous action at  $t_1$ .

If an individual's diachronic responsibility for an individual action (Type 2) was always equivalent to that individual's synchronic responsibility for that individual action (Type 1), then this would be a distinction without a difference. For example, many theorists explicitly or implicitly believe that diachronic responsibility is a straightforward matter of personal identity.<sup>11</sup> According to this view, diachronic responsibility does not come apart from synchronic responsibility so long as personal identity holds; if an individual at  $t_1$  is blameworthy to degree  $d$  for some action that occurs at  $t_1$ , then some individual at  $t_2$  is blameworthy to degree  $d$  for the action that occurred at  $t_1$  if and only if the agent at  $t_2$  is personally identical to the agent at  $t_1$ .

---

<sup>10</sup> Note that the debate in the literature about whether responsibility is essentially historical is a debate most fundamentally about the proper conditions of Type 1 attributions. It is not directly a debate about the conditions of Type 2.

<sup>11</sup> Locke, for example, appears to endorse this view: "In this personal identity is founded all the right and justice of reward and punishment" (1694, 46). Many others have agreed.



I have argued elsewhere that personal identity is neither necessary nor sufficient for individual diachronic responsibility.<sup>12</sup> For present purposes, let me briefly review those arguments. To see that personal identity is not necessary consider the following case. Suppose that Bonnie has robbed a bank at time  $t_1$  and suppose that she is blameworthy for doing so to some degree  $d$ . Now imagine that in order to confuse the police Bonnie freely and knowingly decides to undergo a fission operation. This operation involves removing her brain from her skull, separating the left and right hemispheres, and placing them into two new brainless bodies. Call the person resulting from Bonnie's left brain hemisphere Lonnie, and the person resulting from Bonnie's right brain hemisphere Ronnie. Suppose further that each brain hemisphere is a functional duplicate of the other such that both Lonnie and Ronnie awake from the operation with Bonnie's (quasi-) memories, personality traits, and all other psychological characteristics. Each are, at time  $t_2$ , delighted at the success of the operation and the subsequent police confusion and begin to devise their next nefarious activity. I have only sketched the details of the example, but insofar as we imagine such a case vividly I think it is intuitively compelling to think that both Lonnie and Ronnie at  $t_2$  are blameworthy to degree  $d$  for the robbery that occurred at  $t_1$ . For example, it would seem very natural to think that Lonnie is fully blameworthy for the earlier action in an alternate scenario in which the right brain hemisphere is tossed into the garbage rather than implanted into a brainless body. And if so, there doesn't seem to be good reason to deny that Lonnie is blameworthy in the scenario in which Ronnie is present as well. It is hard to see how a fact extrinsic to Lonnie, such as whether Ronnie exists or not, could affect Lonnie's

---

<sup>12</sup> See Khoury 2013 where I develop the distinction between individual synchronic and diachronic responsibility for individual action in more detail. Also see Shoemaker 2012 and Matheson 2014.

blameworthiness for the earlier action.<sup>13</sup> If both Lonnie and Ronnie are blameworthy to degree  $d$  for the earlier robbery, then personal identity is not necessary for individual diachronic responsibility for individual action. This is because both Lonnie and Ronnie are diachronically responsible for the earlier action but both cannot be personally identical with Bonnie. This is because personal identity is a transitive relation. Thus, if Lonnie is personally identical to Bonnie, and Bonnie is personally identical to Ronnie, then Lonnie is personally identical to Ronnie. But this is impossible; Lonnie and Ronnie are not one and the same. Thus, personal identity is not necessary for individual diachronic responsibility for individual action.<sup>14</sup>

We do not, of course, interact with the products of fission in our daily lives (though perhaps we will at some point in the future). What is more significant to the world with which we are acquainted is the fact that personal identity is not sufficient for individual diachronic responsibility for individual action. Return to the case of Red who committed a murder at time  $t_1$ . Suppose that the Red of our example is Ellis Boyd ‘Red’ Redding (played by Morgan Freeman) from the 1994 film *The Shawshank Redemption*. In the film Red is incarcerated in the Shawshank Penitentiary for an act of murder that he committed when he was a teenager. But the film takes place decades later when Red is an aging man. The kind and gentle character that we are confronted with in the film shows no semblance to any collection of beliefs or values that could lead one to commit murder.

---

<sup>13</sup> As Parfit asks in relation to the question of survival in such cases, “How could a double success be a failure?” (1984, 256).

<sup>14</sup> One might object to this claim by appealing to four-dimensionalism. According to the objection, prior to the fission operation there were actually two distinct but overlapping space-time worms occupying the Bonnie body. While this move might help one resist the fission cases as counterexamples to the claim that personal identity is necessary for diachronic responsibility, one will still be forced to deal with further counterexamples. This is because four-dimensionalism, as such, does not provide an account of the relation that unifies two person-stages as stages of the same space-time worm. The four-dimensionalist must appeal to some account of the unity relation, whether it be psychological or biological, and there are then further counterexamples to the necessity claim given those particular accounts. See Shoemaker 2012.

Rather, Red strikes us as one of the most decent characters in the film. I claim that given what he is like at that later time, Red the old man (at  $t_2$ ) is (at least) less blameworthy for the murder than is Red the teenager (at  $t_1$ ). His blameworthiness has diminished given the psychological changes that he has undergone over that time.<sup>15</sup> If so, then personal identity is not sufficient for individual diachronic responsibility for individual action on the natural assumption that Red at  $t_2$  is personally identical to Red at  $t_1$ .<sup>16</sup>

What, then, is the relation that underlies attributions of individual diachronic responsibility for individual action, if not personal identity? The reason that Lonnie and Ronnie, intuitively, are diachronically blameworthy for the robbery appears to have to do with the fact that their intentional attitudes are relevantly related to the intentional attitudes of Bonnie that gave rise to the robbery. Both Lonnie and Ronnie possess, identify with, and endorse the very same attitudes that led Bonnie to commit the crime. And the reason that Red at  $t_2$  seems less (or not) blameworthy for the murder at  $t_1$  is that his intentional attitudes at  $t_2$  are not relevantly related to the intentional attitudes of Red at  $t_1$  that gave rise to the murder. Rather than being about personal identity, individual diachronic responsibility is, I suggest, a matter of *psychological connectedness*. Two states are psychologically connected when there is both a causal and similarity relation between them. For example, a memory is psychologically connected to some experience when the experience causes the memory (in the right way), and the memory is of that earlier experience (that is, they share intentional content). The same is true of other

---

<sup>15</sup> One might object that Red at  $t_2$  is just as blameworthy for the murder as is Red at  $t_1$ , it is simply that it is less appropriate to overtly blame Red at  $t_2$  than Red at  $t_1$ . I agree that it would be less appropriate to overtly blame him at  $t_2$  than at  $t_1$ , but I believe that the best explanation of this judgment is simply that he is less blameworthy at  $t_2$  than at  $t_1$ . See Khoury 2013, 737-738, 744. Also see Khoury and Matheson ms.

<sup>16</sup> Note that Red at  $t_2$  is personally identical with Red at  $t_1$  on the two most popular accounts of personal identity: the psychological criterion and the biological criterion.

intentional attitudes such as beliefs and desires. A desire at a time is psychologically connected to a desire at a later time to the extent that there is an appropriate causal and content relationship between them. To speak roughly, this is simply the way in which intentional attitudes persist across time.<sup>17</sup>

Psychological connectedness needs to be distinguished from psychological continuity. Psychological continuity is the ancestral relation of psychological connectedness. This means that psychological continuity is made up of overlapping chains of (strong<sup>18</sup>) psychological connectedness. An important implication of this is that psychological continuity is a binary and transitive relation and, for that reason, does not entail the existence of any particular distinctive psychological features, while psychological connectedness is not transitive and does entail the existence of particular distinctive psychological features. For example, I am now psychologically continuous with my 2 year old self but I am not now (very) psychologically connected with my 2 year old self. Myself at 3 years old remembered the experiences of myself at 2 years old, myself at 4 years old remembered the experiences of myself at 3 years old and so on. Because there are these overlapping chains of direct psychological connections (e.g. memory, belief, and desire), I am now psychologically continuous with myself at 2 years old. But there are few, if any, direct psychological connections between my current self

---

<sup>17</sup> See Parfit 1984, 204-207. The qualification is needed because persistence will hold only when connectedness holds uniquely (i.e. one-one, not one-many). But because whether this relation holds uniquely or not is arbitrary, psychological connectedness captures “what matters” in the persistence of psychological states.

<sup>18</sup> Parfit 1984, 206 defines strong psychological connectedness as *enough* connections to ensure that identity holds day to day, on a revised Lockean view of personal identity. Strong connectedness, then, is a non-scalar binary relation unlike psychological connectedness *simpliciter*. Appeal to overlapping chains of strong connectedness (i.e. psychological continuity) is necessary in order to craft an account of personal identity from the notion of psychological connectedness. This is because personal identity is a form of numerical identity and so must have the logic of numerical identity (e.g. it must be non-scalar and transitive), but psychological connectedness does not have the appropriate logic (e.g. it is scalar and intransitive).

and myself at 2 years old. I do not now, for example, remember any experiences of myself at 2 years old.

On the account that I defend, individual diachronic responsibility for individual action is directly sensitive to psychological connectedness with the intentional attitudes that gave rise to the action in question.<sup>19</sup> Suppose that at  $t_1$  some agent commits some bad action because of some criticizable collection of intentional attitudes and is therefore blameworthy to degree  $d$  at  $t_1$  for the action. If the agent at some later time  $t_2$  is maximally psychologically connected to those intentional attitudes that gave rise to the action, then the agent is, at  $t_2$ , blameworthy to that same degree  $d$  for the earlier action. If the agent at  $t_2$  is psychologically connected to the motivations that gave rise to the action at  $t_1$  only to a small degree, then the agent at  $t_2$  is only blameworthy for the earlier action to a small degree. This is, I take it, the case with Red from *The Shawshank Redemption*.

### **Type 3**

Responsibility attributions of Type 3 involve attributing synchronic responsibility to a collective agent for some collective action. For example, to say that British Petroleum at  $t_1$  is blameworthy for acting negligently at  $t_1$  is to make an attribution of Type 3. Various accounts of Type 3 attributions have been put forward. For example, eliminativists claim that talk of collective agency is either in error or else merely shorthand for talk of individual agency.<sup>20</sup> On this conception, presumably, Type 3 attributions are either false or else merely a handy device to refer to a set of Type 1 or Type 2 attributions.

---

<sup>19</sup> See Khoury 2013, and Khoury and Matheson ms.

<sup>20</sup> This language comes from List and Pettit 2011.

Other theorists reject this form of eliminativism. French, for example, has influentially argued that Type 3 attributions are not always reducible to Type 1 or 2 attributions.<sup>21</sup> Rather, a collective can be functionally structured in such a way that grounds attributions of moral responsibility, in part, because that functional organization can ground attributions of intention. Appealing to a Davidsonian conception of action, some event counts as an action only if there is a true description of that event under which it is intentional.<sup>22</sup> French argues that a corporation's internal decision structure (its CID Structure), which consists of a flowchart delineating the hierarchical roles of the individual members and a set of rules that specifies the conditions under which a corporate decision has been made, grounds these corporate intentional descriptions. They ground the identification of the corporate reasons for which the corporation acts and which may not coincide with the individual reasons for which the individual members act.

Christian List and Philip Pettit have recently rigorously defended the idea of collective moral agency along similar functionalist lines.<sup>23</sup> They argue first, that it can make good sense to adopt what Daniel Dennett has called "the intentional stance" towards some collectives. Adopting this stance towards an entity involves attributing intentional attitudes, such as beliefs and desires, towards the entity as a means of explaining its behavior.<sup>24</sup> For example, just as attributing such attitudes can prove indispensable in explaining the behavior of humans, so too can it prove indispensable when explaining the behavior of some collectives. It is relatively uncontroversial that

---

<sup>21</sup> See French 1984, 1998.

<sup>22</sup> See Davidson 1980.

<sup>23</sup> See Pettit 2007 and List and Pettit 2011.

<sup>24</sup> See Dennett 1987.

collectives can be organized in such a way that they can form beliefs about the way the world is, desires concerning how they want the world to be, and that they can act on the basis of those beliefs so as to realize those desires. For example, a given collective may hold a vote to decide whether it, qua collective, will endorse some proposition or seek to realize some goal. None of this is particularly surprising.

What is surprising is that ensuring that a collective has consistent intentional attitudes gives rise to autonomy at the level of the collective.<sup>25</sup> If a collective is to count as an agent it must minimize the possibility that it has inconsistent beliefs and desires, otherwise it could not act so as to realize its desires on the basis of its beliefs thereby undermining the very possibility of agency. In order to guarantee this minimal condition of agency, List and Pettit argue, a minimally complex collective that holds attitudes on logically related propositions must be organized in such a way that gives rise to the possibility that the collective has an attitude that is not shared by *any* of its members. List and Pettit argue further that since collectives can consider and endorse propositions with moral content (because their individual members can bring these propositions up for consideration) and because collectives can then control what they do on the basis of their collectively endorsed attitudes, collectives can meet the conditions of moral agency. Of particular note is the possibility that a collective is morally responsible in a way that none of its members are (say, because the collective endorses a morally objectionable proposition that is not held by any of its members).

The basic idea of accounts of this kind, which are best understood as accounts of Type 3 (collective synchronic responsibility for collective action), is that collectives can count as moral agents because they have and can freely act upon intentional attitudes

---

<sup>25</sup> See Pettit 2007 and List and Pettit 2011.

such as beliefs and desires, and these attitudes can have moral content. A collective's synchronic responsibility for a collective action is determined in large part, just as is an individual's synchronic responsibility for an individual action, by the attitudes expressed by the agent in the action.<sup>26</sup> At this point, I merely want to note my sympathies with accounts in this spirit and discuss the other types of responsibility with this analysis in hand.

#### **Type 4**

A responsibility attribution of Type 4 attributes to a collective diachronic responsibility for a collective action. For example, to say that British Petroleum is now blameworthy for its past negligent actions that led to the Gulf oil spill is to make an attribution of Type 4.

Just as it is rather natural, but mistaken, to think that individual diachronic responsibility for individual action is a matter of personal identity, one might similarly think that collective diachronic responsibility for collective action is a matter of collective identity, understood as a form of numerical identity.<sup>27</sup> On this view, collective identity is both necessary and sufficient for diachronic collective responsibility for collective action. If a collective at  $t_1$  is blameworthy to some degree for some collective action that occurs

---

<sup>26</sup> I take it that one of the primary insights of Strawson's 1962 landmark essay is that this is what moral responsibility is most fundamentally about. Silver 2002, 2005 and Tollefsen 2003 have defended a Strawsonian approach to collective responsibility according to which appropriate collectives can be the proper target of the reactive attitudes. Recently, Björnsson and Hess 2016 have argued that appropriate collectives can also be the proper *bearers* of the reactive attitudes. In my view the Strawsonian approach is complementary to, rather than an alternative to, the functionalist approaches of French and List and Pettit.

<sup>27</sup> Janna Thomson appears to invoke this view when she says: "They [properly organized collectives] ought to act responsibly, and when they do wrong they ought to make recompense. Since these collectives persist through time and, in many cases, through the generations, so, it seems, do their responsibilities" (2006, 158; also see her remarks on 159-160). Abdel-Nour endorses the claim that identity is necessary and sufficient for both individual and (arguably) collective diachronic responsibility: "No matter when the agent's role as cause of a bad state of affairs is discovered or established, this discovery incurs a responsibility on her merely by virtue of her identity with the agent who performed the deed" (2003, 697). Also see Van den Beld 2002.



at  $t_1$ , then some collective at a later time  $t_2$  is blameworthy to that same degree for that very action if and only if the latter is numerically the same collective as the former. If this view were true, then the distinction between Type 3 and Type 4 would be a distinction without a difference. But this view is mistaken. Collective numerical identity is neither necessary nor sufficient for collective diachronic responsibility for collective action.

First, for reasons parallel to the above argument that personal numerical identity is not necessary for individual diachronic responsibility for individual action (Type 2), collective numerical identity is not necessary for collective diachronic responsibility for collective action (Type 4). Imagine that a terrorist organization  $O$  carries out an attack  $X$  at time  $t_1$ . Imagine that the organization is structured in such a way that it meets the conditions of collective synchronic responsibility, such as those outlined in the previous section, and hence, is blameworthy to some degree  $d$  for its attack at time  $t_1$ . Now imagine that in the subsequent confusion and chaos following the attack, the members of the terrorist organization  $O$  are separated into two groups. Call the resulting groups  $O_1$  and  $O_2$ . Each half, suppose, reasonably believes that it is all that remains of  $O$  and that the others have perished. This is, in effect, a case of collective fission. Suppose that at a later time  $t_2$   $O_1$  and  $O_2$  continue on as before, exemplifying the same organizational structure as  $O$ , and pursuing the same goals in accordance with the same policies and so on. Suppose that, at  $t_2$ , both continue to fully reflectively endorse the attack  $X$  which occurred at  $t_1$ . Insofar as we imagine such a case vividly, I think it is intuitively compelling to think that both  $O_1$  and  $O_2$  at  $t_2$  are blameworthy to degree  $d$  for the attack  $X$  which occurred at  $t_1$ . For example, it would seem very natural to hold that  $O_1$  at  $t_2$  is blameworthy for the attack  $X$  in a similar case in which  $O_2$  does in fact perish (if one is

not skeptical, in general, of Type 4 attributions). It is hard to see how a fact extrinsic to O1, such as O2's survival or death, could affect O1's responsibility for X. But if both O1 and O2 are blameworthy for the earlier attack, then collective numerical identity cannot be necessary for collective diachronic responsibility for collective action. Both O1 and O2 are diachronically responsible for the attack but both O1 and O2 cannot be numerically identical with O, for this would imply that O1 and O2 are numerically identical with each other which is absurd; they are not one and the same, they are two independent systems of belief and desire.<sup>28</sup> Hence, identity is not necessary for collective diachronic responsibility for collective action.<sup>29</sup>

And for reasons parallel to the above argument that personal numerical identity is not sufficient for individual diachronic responsibility for individual action (Type 2), similarly collective numerical identity is not sufficient for collective diachronic responsibility for collective action (Type 4). Imagine that at time  $t_1$  a car manufacturer M intentionally acts so as to deceive a national regulatory agency as well as its customers, by installing defeat software in its vehicles that cheats on emissions tests. Suppose that the actions of the car manufacturer ground an attribution of Type 3, collective synchronic responsibility for its collective action A: *M's freely and knowingly cheating on emissions tests*. Now suppose that 50 (or 500 or 5000) years have passed. Imagine that during that time M altered its organizational structure and policies that gave rise to the earlier emissions cheating, made sincere public apologies for doing so, and offered proportional restitution to those affected. Imagine that now, at  $t_2$ , M voluntarily meets environmental regulations much more stringent than those required by law and is widely recognized to

---

<sup>28</sup> This language comes from List and Pettit 2011, 34.

<sup>29</sup> One might appeal to four-dimensionalism in order to attempt to resist this counterexample. The same remarks apply in the collective context as do in the individual context. See note 14.

set the bar for environmentally friendly vehicles in the automotive industry. Suppose that M at  $t_2$  shares none of the intentional attitudes of M at  $t_1$  that gave rise to the past action A. Insofar as we imagine such a case vividly, I think it is intuitively compelling to think that M at  $t_2$  is (at least) less blameworthy for A than is M at  $t_1$ . That is, M's blameworthiness for A has diminished with the passage of time in virtue of the changes it has undergone. If so, then identity is not sufficient for collective diachronic responsibility for collective action (Type 4), on the assumption that M at  $t_2$  is numerically identical with M at  $t_1$ .<sup>30</sup>

I suggest that the same relation that grounds individual diachronic responsibility for individual action (Type 2) grounds collective diachronic responsibility for collective action (Type 4). The reason that O1 and O2 at  $t_2$  are blameworthy for the earlier action committed by O at  $t_1$  is because both O1 and O2 are fully connected to the intentional attitudes that led to the earlier attack. And the reason that M at  $t_2$  is less blameworthy for A than is M at  $t_1$ , is because M at  $t_2$  is no longer connected to the intentional attitudes that gave rise to A. Recall that the relevant notion of psychological connectedness, what

---

<sup>30</sup> The plausibility of this claim depends, in part, on the correct account of collective numerical identity. Suppose that the correct account appeals to, as do the most popular accounts of personal numerical identity, some form of continuity such as psychological or attitudinal continuity, or some form of physical continuity. Continuity is a relation that is non-scalar and transitive, as it must be if it is to play the appropriate role in an account of numerical identity (which is non-scalar and transitive). Though these logical features of the continuity relation allow it to play the appropriate role in an account of numerical identity, they entail the possibility that continuity can hold in the absence of any "direct connections" whatsoever. That is, continuity is a relation that is made up of overlapping chains of some other underlying relation. There can be overlapping chains of that underlying relation over some duration (and so continuity) even if there are no direct chains of that underlying relation over that duration (and so no direct connections). If the correct account of collective numerical identity looks something like, for example, Parfit's Psychological Criterion (1984, 207), then it will be possible that a collective at one time is numerically identical with a collective at another time even though they share no attitudes whatsoever. For example, it would be possible for a collective at  $t_1$  that is an attitudinal and functional duplicate of the Nazi SS to be uniquely continuous with, and for that reason numerically identical with, an attitudinal and functional duplicate of Oxfam at a later time  $t_2$ . It is in light of such cases that that claim that collective numerical identity is sufficient for diachronic responsibility is most difficult to sustain. On this point in the individual context see Khoury and Matheson ms.

might be better referred to in this context simply as attitudinal connectedness, concerns (what matters in<sup>31</sup>) the persistence of intentional attitudes across time. The picture I am defending is this. When an agent acts it does so on the basis of a collection of intentional attitudes, for example beliefs and desires.<sup>32</sup> An agent's synchronic responsibility, whether individual (Type 1) or collective (Type 3), is largely a matter of the moral quality of these attitudes. And an agent's diachronic responsibility, whether individual (Type 2) or collective (Type 4), is largely a matter of (what matters in) the persistence of those attitudes.<sup>33</sup> Attitudinal connectedness is the proper way to understand (what matters in) the persistence of intentional attitudes, hence, attitudinal connectedness is the relation that grounds ascriptions of diachronic responsibility.

Various accounts of Type 4 have been put forward (particularly in the context of national responsibility for historic injustice). I think that the above account commands a great deal of intuitive appeal. But another reason to prefer this account over rivals concerns theoretical parsimony. The fact that we can explain both individual diachronic responsibility for individual action (Type 2) and collective diachronic responsibility for collective action (Type 4) on a single theoretical model is a mark in its favor.

## **Types 5-8**

I will discuss attributions of Types 5-8 together because they share an important feature.

All these attributions involve an asymmetry between subject and object. That is, while

---

<sup>31</sup> See note 17.

<sup>32</sup> I speak of beliefs and desires for the sake of simplicity. We may need to include reference to other intentional attitudes, notably plans and concerns. See Bratman 1987 and Frankfurt 1982.

<sup>33</sup> Notice that this formulation allows for the possibility that there are additional requirements. For example, one might wish to hold that diachronic responsibility requires that the agent in some way freely chooses to continue to have the relevant attitudes.

attributions of Types 1-4 all involve either individual responsibility for individual action (whether synchronic or diachronic), or collective responsibility for collective action (whether synchronic or diachronic), attributions of Types 5-8 involve a contrast between subject and object. They involve either individual responsibility for collective action (whether synchronic or diachronic), or collective responsibility for individual action (whether synchronic or diachronic).

One might attempt to develop an account of Types 5-8 that takes its cue from the above analysis of Types 2 and 4. The notion of attitudinal (or psychological) connectedness, it was argued, grounds ascriptions of Types 2 and 4. An individual is morally responsible for some individual action in the past to the extent that the individual is connected to those intentional attitudes that gave rise to the action. Similarly, a collective is morally responsible for some past collective act to the extent that it is connected to those intentional attitudes that gave rise to the past collective action. (Recall that attitudinal states are connected, in this way, to the extent that there exists a proper causal and similarity relation between them.)

In a similar vein, one might argue that an individual member of a collective is synchronically or diachronically responsible for an act of the collective (Type 5 or 6) to the extent that the individual is connected, in this way, to the attitudes that gave rise to the collective act. Suppose that a racist organization collectively acts to cause harm. On the view being suggested, a member of the organization may be responsible for that very collective act to the extent that she is connected to the attitudes that led to the action.

To put it another way, recall the general account of Type 3, collective synchronic responsibility for collective action, discussed above. Collectives can be organized in such

a way that it makes sense to adopt the intentional stance towards them. We can then explain collective actions by attributing intentional attitudes to the collective such as beliefs and desires. And these intentional attitudes may have moral content that grounds attributions of Type 3 (and [what matters in] the persistence of which grounds attributions of Type 4). That is, collectives can act with a morally evaluable quality of will. The view being currently considered suggests that an individual can be responsible for a collective act to the extent that her individual quality of will is relevantly connected to the collective quality of will. This could occur either synchronically (Type 5) or diachronically (Type 6).

Suppose that an individual member of an appropriately structured racist organization votes in favor of pursuing certain racist policies that the organization then adopts and acts upon. On this view, the individual is responsible for the collective act to the extent that her quality of will is similar and properly causally related to the collective quality of will that gave rise to the collective act.<sup>34</sup> Similarly, an individual may be diachronically responsible for a past collective act to the extent that the individual is connected to the past collective quality of will.<sup>35</sup>

The prime advantage of this view is that it has the potential to offer a subtle account of the way that responsibility for collective acts distributes to individual members. It allows that the responsibility of members is not necessarily on a par, but

---

<sup>34</sup> According to this view, if an individual member votes against that policy of collective action then she may not be responsible for that particular collective action because she is not relevantly connected to the collective's "springs of action." It should be emphasized, however, that this in no way implies that there is nothing else for which the individual is responsible (e.g. her becoming a member).

<sup>35</sup> There are a number of accounts in a similar general spirit. For example, Abdel-Nour 2003 argues that an agent can be properly connected to the past action of another (and arguably of a past collective like a nation) via the emotion of pride, and that this grounds attributions of responsibility. According to Abdel-Nour, "when there is national pride, there is national responsibility" (713). May 1992 argues that one can be responsible for the acts of a distinct agent in virtue of sharing the relevant attitudes that gave rise to the act.

tracks their relation to the morally relevant features of the collective.<sup>36</sup> A similar account might be given for attributions of Types 7 and 8.

This is just a brief sketch in broad strokes and a fully developed account would need to address further questions. For example, in what way is the causal direction relevant? An individual may be attitudinally connected to a collective will because the individual's attitudes are, in part, a cause of the collective will (e.g. by means of voting). Alternatively, an individual may be attitudinally connected to a collective will because the collective's attitudes are a cause of some of the individual's relevant attitudes (e.g. by means of indoctrination). There may be other possible ways in which an individual's attitudes may be related to the collective attitudes in a way that grounds these types of responsibility attribution. For example, one might appeal to the supervenience relation that holds between the collective attitudes and the individual attitudes.<sup>37</sup> Furthermore, it seems that there may be important distinctions to be made amongst the relevant attitudes. Collective action involves sets of collective belief and desire. It may be that some of those beliefs and desires (or some combination of them) are morally innocuous while others are not. Thus, it appears to matter for responsibility to which collective attitudes in

---

<sup>36</sup> Contrast with, for example, David Silver's 2002 account. While he helpfully frames the issue of Type 6 attributions as one concerning an individual's proper ownership of past collective action, his positive account of this ownership seems to imply that there are no differences amongst members with respect to their responsibility for a collective act. This is because, on his account, one owns the past actions of a collective of which one is a member, simply in virtue of the fact that one is a full-fledged member of that collective. But, arguably, it is implausible that all members of a collective are all equally responsible for the past acts of the collective. Some, for example, may reflectively endorse (i.e. be attitudinally connected to) the attitudes that gave rise to the past action to a greater extent than others, and this seems morally relevant.

<sup>37</sup> One can accept that collective attitudes supervene on individual attitudes without accepting that a collective's attitude towards a proposition supervenes on the attitudes of its individual members with respect to that very proposition. See List and Pettit 2011, 64-73. It's worth noting that the relevant supervenience relation may not be a form of attitudinal connectedness, as defined here, insofar as the latter but not the former involves a causal relation.

particular the individual is connected even holding fixed the collective action and the degree of the connection.<sup>38</sup>

While I am inclined to think that an account along the general lines suggested above is the most promising account of the truth of attributions of Types 5-8, I am hesitant to endorse the view that such attributions are possibly true. Notice that all plausible accounts of Types 5-8, including the above, will share a general structural feature, namely, that they make use of what we can call a *tracing strategy*. This strategy involves “tracing” from the object of the responsibility attribution to some distinct property of the subject that is thought to ground the responsibility attribution.<sup>39</sup> For example, on some accounts of Types 5 and 6 an individual can be responsible for a collective act only if we can trace from the collective act to an appropriate participatory intention.<sup>40</sup> Other accounts hold that we must be able to trace back to an appropriate attitude or an act of identification.<sup>41</sup>

All such accounts involve specifying some distinctive basis on which the subject is morally responsible for the object.<sup>42</sup> And on each of these accounts, because the basis and object are ontologically independent the relation between them is contingent such

---

<sup>38</sup> The same issue arises in the case of individual diachronic responsibility for individual action (Type 2).

<sup>39</sup> The notion of tracing is familiar in the literature on free will and moral responsibility in the individual context. Many theorists believe that an appeal to tracing is necessary in order to explain a range of cases of moral responsibility such as drunk driving cases. For example, suppose that an agent freely and knowingly becomes inebriated at a party and subsequently drives home and kills a pedestrian. Stipulate that the agent lacks the control required for moral responsibility at the time of the accident. But, intuitively, the agent is responsible for killing the pedestrian. Many think that we can properly explain the agent’s responsibility only if we trace back from the control-deficient event (the killing of the pedestrian) to some non-control-deficient event (e.g. the agent’s getting drunk). See, for example, Fischer and Ravizza 1998, and Fischer and Tognazinni 2009. Some theorists have raised problems for tracing and offered alternative accounts. See Vargas 2005, Khoury 2012, King 2014, Shabo 2015, and Agule 2016.

<sup>40</sup> See Kutz 2000.

<sup>41</sup> See May 1992 and Abdel Nour 2003.

<sup>42</sup> Kutz 2000, 122 makes explicit the distinction between object and basis.



that the basis and object can come apart.<sup>43</sup> And this feature, in turn, entails a form of moral luck.

Suppose, for example, that the given account holds that an individual subject  $S_i$  is blameworthy for collective act  $O_c$  on the basis of some individual property  $P$  of  $S_i$ . Because  $P$  (whether it be a participatory intention, a faulty attitude, an act of identification, or whatever) is ontologically independent from (and for that reason, only contingently related to)  $O_c$ , we can imagine an alternate scenario in which  $S_i$  has  $P$  but  $O_c$  does not occur. On the given account, there is a difference in  $S_i$ 's blameworthiness across the two scenarios despite the fact that there is, by hypothesis, *no difference* with respect to the basis of  $S_i$ 's blameworthiness across the two scenarios.  $S_i$  is blameworthy for  $O_c$  in the first scenario but not in the second (because it does not occur). But whether  $O_c$  occurs or not is a matter independent of the basis on which  $S_i$  is held responsible. This result is as implausible, by my lights, as the claim that the mere presence of a bird in the path of one's bullet can lessen one's blameworthiness in the context of a murder attempt.

Types 5-8, then, entail a form of moral luck that is relevantly similar to resultant (or consequential or outcome) moral luck.<sup>44</sup> Because the basis and object are

---

<sup>43</sup> This is true with respect to the actual above accounts referenced. Here is a sketch of an argument for the general claim that this relation must be contingent. Note first, that no plausible basis of an individual's responsibility for a collective act, such as her own individual contribution, will itself be sufficient for the occurrence of the collective act. This is because collective action is only possible in the context of the contributions of multiple individuals (i.e. a collective).

Secondly, one can argue that no particular individual contribution is itself necessary for the collective act by appealing to the multiple realizability of collective action (see List and Pettit 2011, ch. 3). French noted early on that certain kinds of collectives, what he called conglomerates, could be multiply realized: "The existence of a conglomerate is compatible with a varying membership. A change in the specific persons associated in a conglomerate does not entail a corresponding change in the identity of the conglomerate" (1998, 44). And it is not implausible to hold that collective action is also multiply realizable such that the actual contribution of any particular member could be fulfilled by another individual. If so, then the actual contribution of any particular member won't be necessary for the occurrence of the collective action. Assessing the extent to which collective action is multiply realizable in this way would require a full account of the individuation of collective action.

<sup>44</sup> See, for example, Williams 1981 and Nagel 1979. Also see Zimmerman 2002.

ontologically independent, facts independent of the basis (such as whether the relevant object occurs) can affect the agent's blameworthiness. Two agents may have exactly the same properties that constitute the relevant basis, but they may be blameworthy for different objects in virtue of facts external to the basis. And this is a form of moral luck. One's blameworthiness can be affected by things other than that in virtue of which the agent is blameworthy. The only way to rule out the possibility of this form of moral luck is to hold that the possible objects of an agent's responsibility cannot be ontologically independent from the relevant basis of responsibility.<sup>45</sup> This, in turn, entails denying the possible truth of Types 5-8 insofar as the object of these responsibility attributions (e.g. a collective act) is ontologically independent of any plausible basis of the attribution (e.g. the individual's own acts).

These anti-moral-luck considerations underlie the claim with which H.D. Lewis began his famous attack on collective responsibility: "If I were asked to put forward an ethical principle which I considered to be especially certain, it would be that no one can be responsible, in the properly ethical sense, for the conduct of another" (1948, 3). Lewis thought this entailed the rejection of collective responsibility altogether because he did not seriously consider the possibility that collectives could be bona fide agents in their own right. But Types 3 and 4, which involve collective responsibility for collective action, do not run afoul of the above principle. Types 5-8, however, do. They all involve an agent's responsibility for the conduct of another agent.<sup>46</sup> For this reason I remain skeptical of the possibility of the truth of such attributions.<sup>47</sup> However, intuitions about

---

<sup>45</sup> I defend this claim, in slightly different terms and in relation to individual responsibility, in Khoury 2012.

<sup>46</sup> Or more carefully, for the conduct of an agent with which the former is not relevantly continuous.

<sup>47</sup> French, forthcoming, has recently raised an objection in a sympathetic spirit against particular accounts of complicity.

moral luck are by no means univocal and I do not expect this objection to gain traction with those not antecedently wary of resultant moral luck. To those who do not see the acceptance of this form of moral luck as a theoretical cost to be avoided, I offer the above account of these forms of responsibility attribution that appeals to attitudinal connectedness.

Before closing I want to briefly consider an objection. I have raised a skeptical worry concerning the possibility of the truth of Types 5-8. These forms of responsibility attribution involve tracing from the object of the attribution to an ontologically independent basis of the attribution. For this reason such attributions give rise to a form of moral luck which, according to the worry, should be rejected. One might object that my own account of Type 2 (individual diachronic responsibility for individual action) and Type 4 (collective diachronic responsibility for collective action) similarly involves tracing and so entails the same form of moral luck. For an attribution of Type 2 or 4 to be grounded we must be able to appropriately trace back from the time to which responsibility is attributed to the time at which the relevant action occurred.

In response, Types 2 and 4 are not relevantly analogous to Types 5-8. With respect to Types 5-8 we must trace back from something that is ontologically independent of the basis (namely, the object) to the relevant basis. It is this structural feature that gives rise to moral luck. But this structural feature is not present in attributions of Types 2 and 4. This is because Types 2 and 4 do not involve tracing back from something that is ontologically independent of the basis to the basis. Rather, these types of responsibility attribution concern, as it were, (what matters in) *the very persistence* of the relevant basis. In this way the objection can be resisted.

In slightly different terms, I accept the principle that we can only be responsible *for* those things that *are* the bearers of responsibility relevant value: in my view, willings that have particular qualities. The most plausible account of types 5-8 involve tracing from something external to the subject's willing to that subject's willing. This entails a form of moral luck. But Types 2 and 4 don't involve tracing back from something external to the willing. They involve (what matters in) the very persistence of the relevant willing.<sup>48</sup>

## **Conclusion**

Peter French's work on individual and collective responsibility has been agenda setting. Building on some of this work, in this chapter I drew attention to three distinctions that arise in the context of individual and collective moral responsibility: subject, object, and temporality. These three binary distinctions give rise to eight types of responsibility attribution. I discussed these various types of responsibility attribution, suggesting that collective responsibility across time can be understood on the model of individual responsibility across time. I considered the possibility that attributions that involve an asymmetry between subject and object can be understood along similar lines, but raised a skeptical concern that such attributions give rise to a form of moral luck. My primary goal, however, has not so much been to defend a particular account of any of these eight types of responsibility attribution, but rather to simply set them out clearly in the hope of gaining some limited degree of clarity in the murky landscape of individual and

---

<sup>48</sup> I should reemphasize, too, that my skepticism only concerns responsibility construed as blameworthiness or praiseworthiness. It is plausible that individuals may have duties or obligations in light of the actions committed by the groups of which they are members. And so it is plausible, *in this distinct sense*, that individuals can be responsible for (i.e. have a duty to respond to) collective harms. See Radzik 2001.

collective responsibility. A landscape which would be all the murkier without French's work.<sup>49</sup>

---

<sup>49</sup> A distant ancestor of this piece was a chapter in my doctoral dissertation of which Peter French was the director. Peter is well known for his writings, but he is also a devoted teacher and mentor and I thank him for his continued support through the years. I also thank Zachary Goldberg for inviting me to contribute to this volume and for comments on the manuscript, and Benjamin Matheson for helpful discussion and comments.

## References

- Abdel-Nour, Farid (2003). "National Responsibility," *Political Theory* 31(5): 693-719.
- Agule, Craig (2016). "Resisting Tracing's Siren Song," *Journal of Ethics & Social Philosophy* 10(1): 1-24.
- Björnsson, Gunnar and Kendy Hess (2016). "Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents," *Philosophy and Phenomenological Research* doi: 10.1111/phpr.12260:1-26.
- Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Davidson, Donald (1980). *Essays on Actions and Events*. Oxford: Oxford University Press.
- Dennett, Daniel (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.
- Fischer, John Martin and Mark Ravizza (1998). *Responsibility and Control: A Theory of Moral Responsibility*. New York: Cambridge University Press.
- Fischer, John Martin and Neal Tognazzini (2009). "The Truth About Tracing," *Noûs* 43(3): 531-556.
- Frankfurt, Harry (1971). "Freedom of the Will and the Concept of a Person," *Journal of Philosophy* 68(1): 5-20.
- Frankfurt, Harry (1982). "The Importance of What We Care About," *Synthese* 53(2): 257-252.
- French, Peter (1976). "Senses of Blame," *Southern Journal of Philosophy* 14(4): 443-452.
- French, Peter (1984). *Collective and Corporate Responsibility*. New York: Columbia University Press.
- French, Peter (1998). *Individual and Collective Responsibility*. Rochester, VT: Shenkman.
- French, Peter (forthcoming). "Complicity: That Moral Monster, Troubling Matters," *Criminal Law and Philosophy*: 1-15.
- Khoury, Andrew (2012). "Responsibility, Tracing, and Consequences," *Canadian Journal of Philosophy* 42(3-4): 187-207.

- Khoury, Andrew (2013). "Synchronic and Diachronic Responsibility," *Philosophical Studies* 165(3): 735-752.
- Khoury, Andrew and Benjamin Matheson. "Is Blameworthiness Forever?" unpublished manuscript.
- King, Matt (2014). "Traction Without Tracing: A Solution for Control-Based Accounts of Moral Responsibility," *European Journal of Philosophy* 22(3): 463-482.
- Lewis, David (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, H.D. (1948). "Collective Responsibility," *Philosophy* 23(84): 3-18.
- List, Christian and Philip Pettit (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Locke, John (1694). "Of Identity and Diversity," in *Essay Concerning Human Understanding*, reprinted in Perry 2008, pp. 33-52.
- Matheson, Benjamin (2014). "Compatibilism and Personal Identity," *Philosophical Studies* 170(2): 317-334.
- May, Larry (1987). *The Morality of Groups*. Notre Dame: University of Notre Dame Press.
- Nagel, Thomas (1979). *Mortal Questions*. New York: Cambridge University Press.
- O'Connor, Timothy (2000). *Persons and Causes: The Metaphysics of Free Will*. New York: Oxford University Press.
- Parfit, Derek (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Perry, John (2008). *Personal Identity*. Berkeley and Los Angeles: University of California Press.
- Pettit, Philip (2007). "Responsibility Incorporated," *Ethics* 117(2): 171-201.
- Radzik, Linda (2001). "Collective Responsibility and Duties to Respond," *Social Theory and Practice* 27(3): 455-471.
- Shabo, Seth (2015). "More Trouble with Tracing," *Erkenntnis* 80(5): 987-1011.
- Shoemaker, David (2012). "Responsibility Without Identity," *Harvard Review of Philosophy* 18(1): 109-132.
- Silver, David (2002). "Collective Responsibility and the Ownership of Actions," *Public*

- Affairs Quarterly* 16(3): 287-304.
- Silver, David (2005). "A Strawsonian Defense of Corporate Moral Responsibility," *American Philosophical Quarterly* 42(4): 279-293.
- Strawson, P.F. (1962). "Freedom and Resentment," *Proceedings of the British Academy*, 48: 1-25.
- Thompson, Janna (2006). "Collective Responsibility for Historic Injustice," *Midwest Studies in Philosophy* 30: 154-167.
- Tollefsen, Deborah (2003). "Participant Reactive Attitudes and Collective Responsibility," *Philosophical Explorations* 6(3): 218-235.
- Van den Beld, Ton (2002). "Can Collective Responsibility for Perpetrated Evil Persist Over Generations?" *Ethical Theory and Moral Practice* 5(2): 181-200.
- Vargas, Manuel (2005). "The Trouble with Tracing," *Midwest Studies in Philosophy* 29(1): 269-291.
- Williams, Bernard (1981). *Moral Luck*. Cambridge: Cambridge University Press.
- Zimmerman, Michael (2002). "Taking Luck Seriously," *Journal of Philosophy* 99(11): 553-576.