2007

Collective Moral Responsibility

Allyson Rudolph
Colby College

COLLECTIVE MORAL RESPONSIBILITY

ALLYSON RUDOLPH
Philosophy Honors Thesis
Prof. Cheshire Calhoun
Fall 2006–Spring 2007

This thesis, entitled “Collective Moral Responsibility,” is submitted in partial fulfillment of the requirements for Honors in Philosophy by Allyson Rudolph on May 15, 2007.

Advisor: Cheshire Calhoun

ACKNOWLEDGEMENTS

For excellence, the presence of others is always required.
~Hannah Arendt

This project would not have been possible without the insight, support, and overall tolerance of a few very important people. First and foremost, thanks to Cheshire. If you hadn't been willing to tell me that my writing "needed a *lot* of work," I would have a very, very long journal entry, rather than a coherent academic argument. Thanks to my family, for sending me the most massive box of sweets on the day of my thesis presentation and for being generally supportive and amazing. Thanks to everyone in the philosophy department, for being so genuinely interested in this project (and for being willing to accept "I have a lot of work to do on my thesis" as grounds for paper extensions). Thanks to my roommates, Mike Deiffenbach, Tom Reznick, Anna Green, Ian Campbell, and Becca Goldstein, and also to my almost-roommates, Kelly Benvenuto, Jack Drury, and Lena Barouh, for putting up with my stressed-out self and not laughing too much when I asked silly questions about truth serum and zombie morality. And thanks to my dogs for hanging out with me in front of the fireplace over January while I highlighted articles, wrote outlines, and grumbled about Enron.

TABLE OF CONTENTS

INTRODUCTION

What is collective moral responsibility? And why should you care? The answer to the former, like any good philosophical question, is largely unresolved. Although writing on collective responsibility has flourished, particularly in the wake of the Holocaust, and despite the existence of an increasingly consistent bibliography of essential writings in the field, there is no definitive authority on the subject. Unlike individual moral responsibility, however, there is little consensus among the ranks. Many theories disclaim the existence or the possibility of collective moral responsibility—a group is just not the sort of thing that can ever be considered a morally responsible agent. No one, they claim, can present a coherent theory of group action that would allow collectives to be admitted into the moral realm.

And that is why you should care about collective responsibility. You interact with collectives every day. You are part of groups, you act cooperatively, and you exist within a community. If it is possible to define "group" in a morally coherent way, wouldn't you want to know?

In this paper, I set out to do just that—define "group" in a way that allows for moral accountability. I begin by looking at moral responsibility broadly, setting out requisites for moral agency. In the first chapter, I will argue that moral responsibility requires causality, awareness, intention, and volition, and that moral responsibility may be meted out in degrees.

Once the basic requirements for assigning moral responsibility are set out, I address three kinds of groups, and attempt to offer models for understanding the moral responsibility of each. I start with institutions—things like businesses and armies—and argue that an institution may be (and ought to be) considered a singular moral agent. From there I move to small groups of individuals united by a shared situation, which I call situational collections. Unable to construe situational collections as singular moral agents, I set out a model for understanding the actions of these groups in terms of shared cooperative activity and shared individual responsibility. Finally, I address issues like racism, in which group members are united by shared attitudes. I present a

model for understanding these shared attitude communities in terms of blame: when blaming a shared attitude community, one is actually assessing the responsibility to the community itself, as well as the individual community members both because of the attitudes they hold and because of their complicity in creating an environment in which material harm or reasonable fear are likely.

I conclude all of these discussions by arguing that there are actually a few models of collective responsibility that allow for groups to operate within the moral realm, but that the real ramification of admitting more members to the moral community is greater responsibility for individuals. Individuals within institutions ought to take responsibility for their own individual actions, and individuals outside of institutions should be vigilant in demanding that immoral institutions change their ways. Members of situational collections must consider their own individual moral responsibilities and work cooperatively to achieve a morally acceptable outcome. Persons who hold attitudes that contribute to harm must take responsibility for their beliefs in radical ways and engage in self-reflection and deep personal change. Each of the models of collective responsibility I present below is in many ways a call for personal reflection on individual interactions with groups and other moral actors.

CHAPTER 1: MORAL RESPONSIBILITY

Can groups (institutions, collections, or random smatterings of individuals) be held morally responsible for actions in the same way that individuals can? One temptation is to say they cannot—the realm of ethics tends to be regarded as a uniquely human endeavor and groups, although made up of people, are not humans. Yet there are many cases in which it is useful and natural to talk about groups as moral agents. As a consumer, I would like to hold the companies I patronize responsible for the moral implications of their business affairs. Ethicists worry about situations in which groups perpetrate actions that could not have been undertaken by individuals alone. These questions lead to an intuitive tug-of-war about the existence of collective responsibility.

There are some arguments claiming that collective responsibility undermines ethics. H.D. Lewis goes so far as to call collective responsibility a "barbarous notion."[i] Lewis, and others who subscribe to his form of individualism, takes it as foundational to ethics that no individual person can be held morally responsible for the actions of another individual.[ii] Collective responsibility, according to Lewis, violates this rule; because collective responsibility goes against a foundational ethical rule, he thinks we must reject the notion of collective responsibility.

The answer to this objection comes from the idea of responsibility distribution. Lewis assumes that at least some cases of collective responsibility assignations entail inappropriate ascriptions of moral responsibility. If individuals are being blamed for the actions of others, or if the moral agency of the group is allowing individuals to escape the yoke of moral agency, Lewis worries that the foundations of ethics are undermined. Answering this objection requires showing that inappropriate claims about individual responsibility do not necessarily occur as a result of assigning collective moral agency.

None of the most compelling descriptions of collective responsibility entail this kind of problem. In the case of institutions, the best account (I think) assigns moral agency to the

corporation itself, and responsibility does not distribute to any individual members of the group. (This, though, is not to say that individuals within a corporation are not responsible for anything involving the corporation, just that they will be responsible for different things.) Similarly, in the case of collectives and random assortments, the responsibility of the group is quite distinct from the responsibilities of the individual group members, and there is no account in which a person is bearing the moral burdens of any other person. Lewis' worry will be addressed in more detail as the full distributability implications of various types of collective responsibility become more apparent.

Another argument against collective responsibility suggests that the notion is incoherent, although not barbarous. It does not make sense to hold a group responsible, because a group cannot possibly satisfy the requirements that we assume for individual responsibility. If a group is not cut out to fulfill the criteria that we use to hold individuals morally accountable, it is not clear how we can assign it moral agency. In order to decide whether or not a group can satisfy the criteria for moral responsibility, the question to ask is: What *is* necessary for moral responsibility?

The first criterion is causality. Talk of responsibility, broadly, revolves around three things: an agent, an action, and an outcome. To hold an agent morally responsible is to blame (or praise) that agent for their actions that have led to an undesirable (or desirable) outcome. For an agent to be morally responsible, the agent's action must be the cause of the outcome in question.

Joel Feinberg is frequently cited for his three-part definition of causality: what he calls "contributory fault."[iii] Feinberg claims that for a person to be considered responsible for certain harms, the person must have acted or failed to act in a way that "made a substantial causal contribution"[iv] to the action in question. Furthermore, the causally connected action must be faulty. Finally, there must be a link between what was faulty and the harmful outcome. As Feinberg says, for the harmful outcome to be someone's fault "the requisite causal connection must have been directly between the faulty aspect of his conduct and the outcome."[v]

One helpful piece of Feinberg's account of moral responsibility for harm is that it explicitly includes acts of omission. It is easy to wonder about such acts when the language of causality comes into play. Can an agent be responsible, for example, for failing to save a drowning child? The clear answer ought to be: yes. The causality requirement should not interfere with holding agents responsible for acts of omission. Because the agent must be causally connected to the *outcome* in question, the nature of the *action* creating that causal link is not at issue. So it does not violate the causality requirement to say that the person who failed to save the drowning child caused the child's death, even though the agent *did* nothing. Feinberg's account is helpful for explicitly incorporating failure to act as a cause.

One challenge to Feinberg's model, though, is this: it is unclear what a "substantial" causal contribution will look like. Qualifiers like "substantial" or "significant" are ambiguous, unnecessary, and morally problematic. It is not possible to say how a significant causal contribution differs from an insignificant one. Explanations of difference between "significant" and "insignificant" will fall victim to arguments from small improvements—it seems intuitively wrong that, on a continuum of significance, the same small improvement that can change something from "very significant" to "very, very significant" can also change something from "significant" to "insignificant."[vi] Any continuum with opposites at either end will lead to worries about vagueness and ambiguity that are very difficult to resolve.[vii] Furthermore, if moral responsibility admits of degrees (which it does, as I will argue later) the "significant" qualifier is unnecessary. As long as there is some level of causality, there will be some degree of responsibility. And finally, when it comes to moral decision-making, our requirements for moral responsibility ought to be fair, but not overly difficult to achieve. With qualifiers like "substantial," it becomes possible to set the standards for moral responsibility too high to be morally useful.

Morally responsible agents, then, must be causal contributors *in some way* to the outcome in question. The first task in refining the notion of causality is to drop the "substantial" qualifier in "substantial causal contribution." In order to hold an agent morally responsible for harm, the

agent must have acted, or failed to act, in a way that caused the outcome in question. There must be an action (including acts of omission) linking the agent to the harmful outcome.

A second challenge to Feinberg's model of responsibility for harm refers to the second section of his account, where he states that the causal conduct must be faulty. What is "faulty"? Feinberg never explicitly states what he means by faulty action, and his article doesn't provide any hints about his meaning. Larry May, who subscribes to Feinberg's model, restates the Feinberg's second section as: "the person's conduct was blameworthy or it was morally faulty in some way."[viii] This seems circular, though. The goal of assigning responsibility is to be able to blame someone for an action that caused harm. To have, as a criterion for causal contribution to harm, that the conduct must be blameworthy, undoes that reasoning. And May's "morally faulty" is no less murky than Feinberg's "faulty."

The emphasis on fault is a problem in another way as well: because of the faulty action requirement, Feinberg's account is only a model of responsibility for *harm*. If the causal contribution must be faulty, then the model can only explain a certain kind of responsibility. However, it is frequently adopted wholesale to explain responsibility in general. As a model of responsibility, Feinberg's causality is misleading; it only works as an account of responsibility *for harm*. So, not only must the "faulty" nature of actions be clarified, it will be useful to identify "faulty's" analogue for moral responsibility for *good* (as opposed to harm).

The task, then, is to explicate "faulty." Rather than leaving it undefined, or defining it in a way that incorporates blameworthiness, a faulty action, for the purposes of assigning moral responsibility for harm, is one that violates some moral norm.[ix] The case of moral responsibility for good is more complicated. We don't praise people for refraining from violating moral norms —it would seem odd to say "I'm so proud of you for *not* committing a morally reprehensible action today!" Moral praise is reserved for the morally exemplary, so perhaps the analogue to faulty action is exemplary action. According to this account, an action must violate a moral norm to make an agent morally responsible for harm, and must exceed a moral norm to make an agent morally responsible for good.

So, the first necessary element of moral responsibility is causality. Causality involves an agent, an action (with "failure to act" understood as an action of omission), and an outcome. In order to blame/praise an agent for an outcome, the agent must have committed an action that violates/exceeds moral norms, and that defiance of moral norms must be the link between the agent and the outcome in question.[x] This model of causality preserves the basic ideas of Feinberg's oft-cited account, and hopefully clears up any ambiguity or confusion.

The second criterion for moral responsibility is awareness. Virginia Held makes an important distinction that takes us from mere causal responsibility to a more complicated moral responsibility—awareness is at the heart of the distinction. She gives the following example: a person throws a bomb through the window of the house, killing a child. That person is both causally responsible for the explosion and morally responsible for the death of the child. But if that same person flips on a light switch that, unbeknownst to her, triggers a bomb that kills a child, she is only causally responsible for the explosion. She is not morally responsible for the death of the child, because she was unaware of the moral ramifications of her action.[xi] Held's epistemic requirements are quite lenient, a point she acknowledges readily: "These requirements are not meant to be equivalent to requirements for complete knowledge…they require that [the agent] be aware of the kind of action he is performing, but not that he know everything about either its empirical or its moral aspects."[xii] (I am using "awareness" instead of "knowledge" precisely because the epistemic requirements are so different from the requirements for complete knowledge.) Some awareness requirement is important for moral responsibility; it seems unfair to hold an agent morally responsible for an action with moral consequences of which she was completely unaware.

Furthermore, agents should be considered morally responsible for failing to acquire awareness of the moral consequences of their actions. It is wrong to claim that a businessperson selling light arms to dodgy characters in war-torn countries should be considered causally, but not morally, responsible for the deaths inflicted by those guns simply because she didn't know what the guns would be used to accomplish. Responsibility, in this case, requires making an attempt to

become aware—in other words, she *should have known better*. She did not immediately understand the moral consequences of her actions, but she is still morally responsible for failing to learn more about a situation that a reasonable person would find potentially laden with moral repercussions. The second requirement for moral responsibility, then, is a certain kind of awareness. A person is morally responsible for her actions only if she a) is aware of the moral ramifications of her actions, or b) could have been reasonably expected to gain such awareness, and did not.

The next criterion for moral responsibility is intention. Intent is central to many discussions of responsibility—there are some who think that agents are just responsible for outcomes they intended.[xiii] But the primacy of intention might be misleading. As Peter French points out, intentions are referentially opaque.[xiv] Whether or not a person intended an event depends upon how the event is described. French uses the death of Polonius in *Hamlet* to show what referential opacity means. Hamlet's killing Polonius and Hamlet's killing the man behind the arras in Gertrude's room are the same action with the same outcome—a dead Polonius. It is true that Hamlet intended to kill the man behind the arras. It is not true that Hamlet intended to kill Polonius—he thought it was the fratricidal Claudius hiding in Gertrude's room. Regardless of description, though, Hamlet is responsible for felling Polonius.[xv] Does the Hamlet example compel us to remove "intention" from the list of necessary criteria for moral responsibility?

Ignoring intention runs counter to most thought about responsibility—intentional agency is a vital part of the debate about moral agency. Fortunately, French has an answer that saves intention: we have to intend *something*. Intention is helpful because it separates actions from events, and only actions will be subject to responsibility judgments.[xvi] So we have to intend the action on some, but not all, descriptions of the action.

French's account of intention is constrained by the causality criterion in a subtle way. The model of causality described above requires that a moral agent act in a way that defies moral norms. French's model is useful for separating actions from events, but makes no distinctions about norm violations. In order for an agent to be held morally responsible, then, there must be at

least one description of the event in which the agent intended the action that defies the moral norm for which we want to hold the agent responsible. Returning to the Hamlet example, imagine a different (and wholly inaccurate) description of the death of Polonius. Say Hamlet heard a rustling behind the arras and assumed there was a zombie hiding there, about to attack Gertrude. Hamlet lunges his sword at what he thinks is the undead and kills Polonius, because of course there is no zombie. Intending to kill a zombie is a very different thing than intending to kill a not-undead human, and zombie-killing does not violate any moral norms that I am aware of. In this situation it makes sense (and seems intuitively correct) to say that Hamlet is causally responsible for the death of Polonius, but that he is not morally responsible.

The third criterion for moral causality, then, is intention. The agent must intend, on at least one description of the event, the moral-norm-defying action that caused the outcome in question. If we are holding the agent responsible for violating a moral norm, then the agent must have intended to violate that norm.

The next criterion for moral responsibility is volition. Volition must be a criterion on its own because it is different than intent. As J.L. Austin points out, intending an action does not imply that the action was performed on purpose, or deliberately. He gives many examples of cases in which an action is intentional, but not "on purpose" or deliberate.[xvii] Although he does not explicitly discuss the difference between volition and intention, his ideas can be extended to include actions that are intended, but not voluntary.

How can an agent act involuntarily? Choice is critical to volition: any act involving a choice is voluntary. Aristotle says that voluntary actions have their origin in the agent, and that "when the origin of the action is in him, it is also up to him to do them or not do them."[xviii] It is not easy to come up with examples of involuntary action, in this case. Aristotle writes that forced actions are involuntary, but coerced actions (although made up of both force and choice) are "taken as a whole" voluntary.[xix] The best candidates for involuntary actions, then, seem to be actions over which an agent has no control. Sneezes, seizures, and fainting spells cannot be rightly said to be voluntary—but they also do not count as actions, because they are not intended

upon any description. An involuntary action, then, would have to be an action that is intended, but not chosen or coerced.

Instinctive responses might be good candidates for truly involuntary action. Pulling your hand away from a hot stove is not an act that is chosen, but it is intended. Actions that come about as a result of choices, but themselves do not allow for choice, might also be instances of involuntary action. Suppose a mountain climber finds herself at a passage too narrow to fit through with all of her gear. She must choose whether or not to go through the passage, but once she has made the decision to take the narrow route, she *must* leave some of her gear behind. The act of dumping the gear is intentional, but not voluntary; the only voluntary action was the decision to take the narrow passage. Once that decision was made, the climber had no choice about whether or not to drop the gear. And cases of actual psychological manipulation also seem to be instances of truly involuntary action. For example, I have been given truth serum and believe that I must answer all questions honestly. In response to a question, I supply an answer that implicates my brother, who I love dearly and would never hurt, in some terrible crime. In this case, I intended to answer the question, because I believed I had to, but the action was not voluntary, because (given the nature of truth serum) I could not have been reasonably expected to hold my tongue. These are the types of truly involuntary actions that do not involve making a choice that are exempt from ascriptions of moral responsibility.

There is an important difference between acting voluntarily and being voluntary with respect to an outcome: A person acting voluntarily may not have been bringing about an outcome voluntarily. What does it mean, then, to be voluntary or involuntary with respect to an outcome?

It will be easier to explain how an agent can be involuntary with respect to an outcome after explaining what it means to be voluntary with respect to an outcome (as opposed to acting voluntarily). The Catholic doctrine of double effect, by which it is permissible to bring about some outcome that is forseen, but not intended, leads to many examples of *outcomes* that are voluntary, regardless of the nature of the action. For instance, the doctrine of double effect is frequently invoked in discussions of terrorist vs. targeted bombing. In both cases, the act of

bombing in voluntary. In the case of terror bombing, the intended outcome is to kill civilians. In the case of targeted military bombing, the intended outcome is to wipe out some strategic target. The death of civilians is a forseen consequence—the bombers killed civilians voluntarily (they could have been reasonably expected to question their commands if they believed they were being instructed to violate an international law), but the intent was not to kill innocents. Regardless of whether or not collateral damage *is* in fact morally permissible, the doctrine of double effect leads to examples in which outcomes are voluntary, regardless of the intent or volition of the causally linked action. To be voluntary with respect to an outcome, then, is to choose to allow some foreseeable outcome.

So, an agent who is involuntary with respect to an outcome is someone who has no control over that outcome. Say I choose to put off shoveling my snowy sidewalk for the ten minutes it will take to make a cup of coffee. I have acted voluntarily in choosing to wait until properly caffeinated to tackle the snow. If, in that ten minutes, my neighbor slips on my icy walkway and breaks her arm, I am not voluntary with respect to that outcome. It is an outcome of my previous voluntary action—I could have chosen to act differently by shoveling the sidewalk before succumbing to my coffeemaker. I could not have chosen a different outcome, however. In this situation, once the decision to forgo shoveling has been made, I am no longer in control of the resulting outcomes. If an agent is not voluntary with respect to an outcome, that agent is not morally responsible for that outcome.

One way that an agent can be involuntary with respect to an outcome is to be unaware of the consequences of the action, as is the case with the snow-shoveling example. But being involuntary with respect to an outcome is not the same as being unaware of the moral consequences of an action. There are ways of being involuntary with respect to an outcome while being aware of the outcome—people are often aware of outcomes over which they have no control. For example, I am aware that every time I take a shower I am using many gallons of water, a nonrenewable resource. Wasting valuable and usable water is a direct consequence of my action, but the depletion of the world's clean water is not something that I can control. I can

change my habits to minimize my environmental impact, but the outcome of my bathing habits is not something that I can choose.[xx] It is important to keep actions and outcomes clear—in the case of environmental degradation, an agent may be morally responsible for failing to act in a more environmentally friendly fashion, but cannot be morally responsible for the foreseeable but involuntary depletion of environmental resources.

Volition, then, is distinct from intention and awareness. An agent is acting involuntarily when her action is not something about which she could have made a choice. Involuntary actions are not subject to assignations of moral responsibility. An agent is involuntary with respect to an outcome when she had no control over that outcome. Similar to involuntary action, an agent cannot be held morally responsible for an outcome over which she had no control (although her actions may be evaluated separately from their consequences).

The above discussion of individual moral responsibility yields four criteria for holding an agent morally responsible. First, the agent's action must have relevantly caused the outcome in question. Second, the agent must have had some awareness of the moral consequences of her actions. Third, the agent must have intended the action on at least one true description of the event. And fourth, the agent must have acted voluntarily (based on a rather easygoing definition of volition). Importantly, each of these criteria can be satisfied to a greater or lesser degree.

It sounds odd to talk about degrees of causality, but much less odd to say that someone directly or indirectly caused something. The notion of indirect causality, though, is a tacit acknowledgement of degrees of causality. Take, for example, a hate crime. Suppose person A severely beats racial minority B for no other reason than B's race. To say that A directly caused the brutality is to acknowledge that A was causally responsible to a very large degree. Suppose further that A might have been apprehended by the police, but just as the cop car was rounding the bend, some close friends of A created a distraction and prevented the cops from witnessing the event. The friends caused the hate crime in that they took an action to prevent police intervention, but their participation was indirect, since they were not actually part of the beating. To say that they indirectly caused the beating is just to realize that they were causally connected

to a lesser degree than the person who actually perpetrated the crime. Further, imagine this beating took place in a town well known for being intolerant of racial minorities, with a history of protecting people like A. The people of the town, then, are still causally connected to B's beating, because their actions contributed to an environment in which racially motivated brutality could easily happen. Naturally, though, they did not cause B's beating to the same extent that A caused B's beating. The people of the town were causally connected to an even lesser degree than A or A's friends. This example shows that, while it might not be natural to use the vocabulary of degrees of causality, there is conceptual space for the idea that causal connections admit of degrees.

Similarly, it is difficult to see how intention might admit degrees. Returning to Hamlet and the death of Polonius, however, shows how different actions can be intended to a greater or lesser degree. Hamlet intended to kill, so the act of killing Polonius was an intentional action. But the act of killing Polonius would have been intentional to a greater degree if Hamlet had intended to kill Polonius. Because Hamlet intended to kill, but did not intend his victim to be Polonius, he intended the act of killing Polonius to a lesser degree than if he had actually intended to kill Polonius. There are more obvious examples of degrees of intention—"I didn't mean to" is an oft-abused appeal to leniency. That agents bring up their intentions in the hopes that less intentionality will lead to less stern judgment shows that, although it may not be natural to talk about degrees of intention, the idea is not uncommon.

It is more natural to talk about degrees of awareness than degrees of causality or intention. There will always be variability as to how much anyone knows about any situation, so of course there will be variability as to how much we know about the moral implications of our actions. Degrees of awareness are incorporated into the U.S. legal system. A child who commits a crime is typically tried as a minor, and will be sent to a juvenile detention center for any conviction. The intuition behind juvenile detention is that children of a certain age are less aware of their actions and of the law than are adults. Children can be tried as adults, however, if they are deemed aware enough of their actions. A six year old who fires a gun is doing so with a much

more limited degree of awareness than a fourteen year old. The intuition behind juvenile detention is the same intuition that allows talk of degrees of awareness.

Volition, too, can be understood in degrees. Some choices are intuitively less voluntary than simple ones. For example, Aristotle classifies coerced actions as voluntary, because, while it is undeniable that the agent is making a choice, it is also clear that the agent would never make such a choice in a situation that was not coerced.[xxi] Cases of duress might lead to unthinkably difficult decisions, but as long as the agent is making a choice, the action is voluntary. Thinking in terms of degrees of volition prevents assigning the same level of volition to a person who commits an act under duress as to a person who acts in relative freedom.

Similarly, there are degrees of volition with respect to outcomes, although the language is not natural. Regardless of the vocabulary, however, it is intuitive to realize that certain outcomes are accepted grudgingly, while others are accepted with less trepidation. It is one thing to allow that, in the course of war, soldiers will die. It is a different thing to allow that, in the course of a war, civilians will die. Both are outcomes, and there are agents who are voluntary with respect to both outcomes, but there is an understanding that, unlike a civilian, the combatant signed up for the war with the knowledge that she might die. Collateral damage seems less voluntary than the combat death of soldiers. Although the language of degrees is awkward here, the intuition that volition comes in degrees remains.

From here it should be clear that moral responsibility itself must come in degrees. If each of the criteria can be satisfied to a greater or lesser extent, then an agent is responsible for their actions just to the extent that they satisfied each criterion. This is not an attempt to turn moral responsibility into a mathematical formula that adds and subtracts and results in a percentage or a score. "Degree" is a misleadingly precise term, but severe precision isn't necessary to grasp the idea an agent can be morally responsible for an action to a greater or lesser extent, depending upon how directly their action caused the outcome, how much they knew about the moral consequences, and how voluntary and how intentional their actions were.

There are at least two objections to the idea that responsibility admits of degrees. First, it is unclear that ordinary talk about responsibility accommodates degrees. When we assign responsibility, we say that someone is or is not responsible. There is no vocabularistic middle ground. Although we don't have precise language for it, I think that everyday talk about responsibility does allow the idea that moral responsibility comes in degrees. It is certainly within the conceptual range of most humans to form a thought similar to "well, she *is* responsible, but you have to take into account the fact that she was under a great deal of pressure. I probably would have done the same." There may not be specific words for "responsible to a great degree" or "responsible, but not in such a way that warrants a lavishing of praise," but everyday talk about moral responsibility does frequently take into account factors that would lessen or amplify the degree of responsibility.

The second argument against degrees of responsibility is that there are no corresponding degrees of punishment or reward in the moral realm. Agents are either deserving of praise, or not deserving of praise. Agents can be blamed, or they can not be blamed. Again, I think that, although the precise language may be lacking, praise and blame are routinely ascribed to a greater or lesser extent. For example, suppose I rescue a cat in a tree, believing it to be my own, only to find out upon rescuing it that the cat is the similar-looking one belonging to my next-door neighbor—a person I happen to despise. If I return the cat grumbling "if I'd known it was *your* cat and not mine, I'd have left it there," I will probably receive less praise than if I had rescued the cat knowing full well its provenance. Similarly, if a person with a gun to her head commits a morally reprehensible action, she will undoubtedly be blamed to a lesser degree than someone who had committed the same action without such compulsion.

So far, I have asked whether or not groups can be held morally responsible for their action. I rejected H.D. Lewis' "no" answer to that question, and asked what it means for *anything* to be morally responsible. A clear understanding of moral responsibility will be useful in responding to the argument that collective responsibility is an incoherent notion. After positing

four criteria for moral responsibility, I argued that moral responsibility must be thought of in degrees. These two notions will allow me (finally) to answer the question "can a group be held morally responsible for its actions?"

One last question of categorization remains: what is a group? The literature on more-than-individual responsibility yields a surfeit of designations: corporations, collectives, conglomerates, aggregates, social groups, random collections—each with its own definition, and all with subtle differences. For this paper, I will define a collection as a group of individuals with something in common (e.g. the collection of brown-eyed Mainers is the group of people who share two traits: having brown eyes and living in Maine). I will focus on three kinds of groups: institutions—collections of people characterized by organizational structures and internal decision-making practices; situational collections—groups of people united by a shared situation; and shared attitude communities—groups in which each member shares some attitude in common.

---

[i] H. D. Lewis, "Collective Responsibility," *Philosophy: The Journal of the Royal Institute of Philosophy* 23 (January, 1948).

[ii] ibid.

[iii] Joel Feinberg, "Collective Responsibility," *Journal of Philosophy* 65 (November 1968).

[iv] ibid.

[v] ibid.

[vi] For a more thorough version of this argument, see Ruth Chang, "The Possibility of Parity," *Ethics: An International Journal of Social* 112, no. 4, pp. 659–688 (2002).

[vii] Again, more on the problems with scales and chaining can be found in Chang (2002), and also John Broome, "Is Incommensurability Vagueness?" in *Incommensurability, Incomparability and Practical Reason*, ed. Ruth Chang (Cambridge, Mass.: Harvard University Press, 1997).

[viii] Larry May, *Sharing Responsibility* (Chicago: University of Chicago Press, 1992).

[ix] Cheshire Calhoun, personal conversation, November 2006.

[x] From this point on, I will mostly be referring to responsibility for harm.

[xi] Virginia Held, "Can a Random Collection of Individuals be Morally Responsible?" *Journal of Philosophy* 67 (July, 1970).

[xii] ibid.

[xiii] J.L. Mackie is one such philosopher. For more, see J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin Books, 1977).

[xiv] Peter A. French, "Principles of Responsibility, Shame, and the Corporation" in *Shame, Responsibility, and the Corporation*, ed. Hugh Curtler (New York: Haven Publications, 1986), 19. Also, chapter three of Peter A. French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984), 215.

[xv] French, *Collective and Corporate Responsibility.*

[xvi] French, *Principles of Responsibility, Shame, and the Corporation.*

[xvii] J. L. Austin, "Three Ways of Spilling Ink," *Philosophical Review* 75 (October 1966), 427–440.

[xviii] Aristotle, ed. Terence Irwin, *Nicomachean Ethics* (Indianapolis, Ind.: Hackett Pub. Co., 1985), 210.

[xix] ibid.

[xx] This is likely true of many cases of environmental degradation—people are increasingly aware of their environmental "footprints," but having the ability to change the actions that lead to negative environmental outcomes is not the same as being able to change the outcomes themselves.

[xxi] ibid.

CHAPTER 2: INSTITUTIONS

Businesses, the military, and the Boy Scouts are all examples of institutions. The burgeoning corporate social responsibility field, the academic pursuit of just war theory, and criticisms of the Boy Scouts' anti-gay policies are all examples of the desire to hold institutions morally responsible for their actions. What is an institution, and can an institution really bear moral responsibility?

I suggest that institutions are *potential* moral agents, and that they ought to become *actual* moral agents. By defining an institution as a group of people characterized by an organizational structure and an internal decision-making process, an ability to accommodate varying membership and varying levels of support, and institutional goals, histories, and basic principles, I hope to address questions about the possibility of institutional moral agency and argue that institutions both can and should be considered moral agents.

The most important feature distinguishing institutions from other collections of individuals is the presence of an organizational structure and an internal decision-making process. Peter French and Virginia Held both suggest that the existence of a decision-making procedure is crucial for distinguishing between kinds of groups. For Held, a group with a decision method is an organized group or a collectivity, as opposed to a random collection of individuals.[xxii] For French, internal decision-making structures distinguish corporations from other groups or collectives. His corporate internal decision-making (CID) structures have two components: a flowchart, and rules.[xxiii] The flowchart refers to the corporation's organization and delineates levels of power and responsibility. The rules are steps that a corporation must take before making a decision in order for that decision to be recognized as one made by the corporation. The CID structure, "when operative and properly activated…accomplishes a subordination and synthesis of the intentions and acts of various biological persons into a corporate decision."[xxiv]

Held's scope is broader than mine—for her, a collectivity is anything with a decision method. I argue that institutions require decision methods as part of a constellation of

characteristics distinguishing them from other groups. French's scope is, I think, narrower than mine—"corporation" seem to designate businesses or groups that are incorporated in the legal sense. Despite differences in scope, French's description of CID structures provides a useful starting place to examine two crucial features of institutions.

French's CID structure can be expanded to describe an institutional decision-making process. Like the organizational structure, an institution's decision process distinguishes it from a non-institutional collection, and is generally obvious and codified. Like French's CID structure, an institutional decision process serves to collate the thoughts and feelings of an institution's individual members and compile them into one institutional decision. Unlike French's account, institutional decision procedures may have varying levels of rigidity, formalization, or standardization. Ultimately, though, they all take the opinions of individuals and depersonalize them to create a single institutional decision that cannot be associated with the particular beliefs or intentions of any individual member. Every institution has some sort of decision-making process, and it is from this unique trait that the argument for institutional moral agency derives its force.

French's idea of a corporate flowchart can be similarly expanded to describe a more general organizational structure. Such a structure plugs members into particular positions and delineates relationships of power based on those positions. The organizational structure of an institution is generally obvious: there are bosses and interns, generals and privates, troop leaders and scouts. Typically, individual members have specifically designated responsibilities and functions within the group. That an institution has, by definition, an organizational structure will be helpful for explaining how and why an institution may be held morally responsible.

So, an institution must have some method by which it can make decisions that are understood as institutional, and it must have an organizational structure. It must also be compatible with varying membership, making it what French calls "conglomerate," rather than "aggregate."[xxv] An institution persists through membership change: a business is the same business, regardless of personnel change; a nation's military is one military, despite casualties,

resignations, discharges, or recruitments; the Boy Scouts are the Boy Scouts, even though scouts may move on and leaders may leave. Even small groups, like sports teams, college clubs, or fraternity chapters, can remain the same institution through changes in membership. Accordingly, the identity of an institution "is not exhausted by the conjunction of the identities of the persons in the organization."[xxvi]

The fact of varying membership is the foundation of the very question of institutional moral responsibility: if an institution were just identical to the conjunction of its individual members, there would be no worry about institutional moral responsibility, because each individual member would be responsible for her own actions. But given that the institution is a singular entity that persists through varying membership, comprising not just its individual members, but also its organizational structure, decision-making method, history, goals, and basic principles, there is a strong intuitive pull to think about institutions as moral agents. Institutional continuity in the presence of changing membership is at the core of questions about institutional moral responsibility.

Institutions have decision methods, organizational structure, and can accommodate varying membership; they also have goals, histories, and underlying principles. Institutions have processes that allow them to make decisions, and decisions are typically made in the hopes of accomplishing some goal. A business might make a decision in order to achieve its goal of turning a profit, or a military may decide on a course of action with the goal of keeping a state safe. Institutional decision-making implies institutional goals. Futhermore, because both an organizational structure and a decision-making process generally involve recordkeeping and transparent communication, an institution will have a history. Institutional histories might be quite formal: there are historians and archivists who research businesses, militaries, governments, the Freemasons, etc. They may also be anecdotal or informal: I am acutely aware of the histories of many of the clubs I have participated in based on the stories of older members and professors, photo albums, and newspaper articles. Finally, institutions have basic principles guiding their membership choices, their organizational structures, their decision-making processes, and their

goals. For many institutions, these basic principles are made obvious through mission statements, codes of conduct, or statements of purpose. For less formal institutions, basic principles may be reflected in a certain method of admitting new members, or a particular type of organizational structure, or an unspoken culture within the group.

The above definition delineates one kind of collective: the institution. A collective counts as an institution if and only if it has an internal decision-making process, an organizational structure, and the ability to accommodate changing membership. An institution will also have goals, a history, and basic principles. All of these defining characteristics can be used to show how an institution can be considered a potential moral agent.

A worry was presented earlier that collectives might not be able to satisfy the conditions that individuals must meet in order to be considered morally responsible agents; in short, that collectives might not be compatible with moral agency. Specifically, are collectives the kinds of things that can act causally, and with awareness, intent, and volition?

Peter French argues that at least one type of collective, the corporation, is actually a moral person. Not only are corporations capable of bearing moral responsibility, they are entitled to the same moral rights as human persons. French advocates a "theory that allows treatment of corporations as full-fledged members of the moral community of equal standing with the traditionally acknowledged residents: human beings."[xxvii] But French's assertion that corporations can be considered moral persons is disconcerting to many philosophers wishing to keep the world of flesh-and-blood humans distinct from non-human metaphysical entities. Objections that corporations cannot expect to be granted health care, security, or even the right to life are abundant.[xxviii] I suggest that by focusing on the institution as a possible moral agent, without the objectionable rights implications of moral personhood, it is possible to understand at least one type of collective as compatible with moral agency. To do this, I will show how an institution can cause, be aware, intend, and act voluntarily by appealing to characteristics defined earlier.

My earlier definition of causality claimed that for an agent to have acted causally to commit some harm, the agent must have committed an action that violates moral norms. An institution can act (or fail to act) causally; that is, can serve as an agent and can commit an action that violates moral norms and that can be causally linked to morally questionable outcomes. An institution can act because it is designed to act. The decision-making process leads to an institutional decision to take some action, and it is just as possible for an institution to act in a faulty way as it is for an individual.

For example, take a state's defense decisions. No one person, even in the most authoritarian of states, develops new defense technology. A complicated system of departments, agencies, and industries are involved in the development of the kinds of weaponry that gets used in contemporary combat operations.[xxix] The defense decisions that emerge from a state's decision-making process are institutional decisions that lead to institutional actions: the *decision* to develop herbicides for strategic use in the Vietnam War led to the *acts* of creating Agent Orange and spraying the toxin over vast swaths of forest in combat areas. The realities of institutional decision-making processes entail institutional action.

Moreover, an institution can act in a morally faulty manner. Institutions, for example, can lie—an action typically considered morally faulty. For example, the inescapable Enron scandal of 2001 was preceded by years of institutional dissembling. The company, as a company, lied to stock analysts to push stock prices up, convincing the analysts that Enron was a moneymaking company in order to achieve "buy" or "strong buy" recommendations. The company explicitly lied to investors, in particular about a bandwidth trading market, to convince the public that Enron was developing feasible, profit-earning projects where there were none. In the years before going bankrupt, Enron as an institution was committing the morally faulty act of lying: to the analysts, the investors, the media, and even the federal government.[xxx]

What about failure to act? Because institutional causality is so heavily based on the existence of a decision-making procedure, failure to act must be construed as an actual *decision* to do nothing. Take, for example, the Rwandan genocide of 1994. United Nations officials had many

indications from Romeo Dallaire, the UN general in charge of peacekeeping in Rwanda, that a massacre was being planned, but made an institutional decision not to make the requested revisions to the mandate that would have allowed Dallaire's troops greater freedom to use force to quell an upsurge of violence. Once the genocide began, most Western nations (including the United States) also made institutional decisions not to commit any troops to stanch the bloodshed, choosing to make the evacuation of their own citizens the extent of their participation in the conflict. The United Nations and countries like the United States refrained, as institutions, from acting, and the outcome of the willful failure to act was a prolonged period of violence and hundreds of thousands of deaths.

An institution can also satisfy the awareness requirement for moral responsibility. According to this requirement, the agent must be either aware of the moral ramifications of action, or in a position to have been reasonably expected to seek such awareness. Institutions can be aware because their organizational structures and decision-making procedures are designed to create institutional awareness and produce decisions based on that knowledge. The goal of an internal decision method is to collect bits of knowledge from across the organizational structure and use the collective awareness to make a decision. The prominence of intelligence agencies in military decision-making is one example of the way institutional awareness comes about through organization and decision-making structures. In practice, the institution may or may not possess *moral* awareness, but every institution has the *ability* to be aware of the moral implications of its actions.

There are certainly cases where institutions act immorally despite acute moral awareness. For example, the Truman administration ordered the bombing of Nagasaki knowing both the scope of devastation from the first atom bomb on the city of Hiroshima and the immanence of Japanese surrender. There are also cases of institutions that change course upon gaining awareness of the moral consequences of their actions. After being presented with evidence that innocent people were being sentenced to death in Illinois courts, the state imposed a moratorium on executions and commuted the sentences of all its death-row inmates to life in prison. Not only

can institutions be aware of the moral consequences of their actions, they can be unaware, and they can gain awareness.

Moving on to intention: can an institution intend an action, and intend to violate a moral norm (or intend an action whose foreseeable outcome violates a moral norm)? Again, the institutional decision-making process is critical to any description of institutional intentionality. The decision-making process is designed to create institutional intent out of individual intentions. Furthermore, the presence of institutional goals and basic principles point to the necessity of institutional intent. Continuing with the military examples, consider Hitler's Germany. It was the stated intent of the Third Reich to eliminate the European Jewish population. The building of concentration camps, the rounding up of Jewish citizens, the listmaking, and the recordkeeping were all institutional actions, developed in response to institutional decisions, based on the institutions' basic principles and goals, that were intended to violate a moral norm, and that resulted in a morally blameworthy outcome. Because of the clarity of the institutional goals and principles, it is easy to see the institutional intent behind the actions.

Finally, can institutions act voluntarily? Agents act voluntarily when the action is something about which the agent could have made a choice. Institutions can act voluntarily because the entire organizational structure and decision-making process is designed to ensure institutional control over choices. By definition, institutions are capable of institutionally voluntary action.

At this point, it seems clear that, by defining institutions as collectives with decision-making methods, organizational structures, varying membership, and institutional goals, histories, and principles, an institution *may* be a moral agent. That is, the means by which an institution acts do not preclude it from bearing moral responsibility—there are no metaphysical barriers keeping institutions from moral agency. Whether or not institutions *are* in fact moral agents is a more delicate question demanding a more normative line of thinking.

It seems to me that, although the possibility of moral agency is an ontological fact, the actual designation of moral agency has to do with normative societal values. The flowers on my

desk are not capable of moral agency—they simply do not possess the required characteristics. People, on the other hand, are generally capable of bearing moral responsibility thanks to certain shared traits that allow us to fulfill the criteria for moral agency. Whether or not something that is capable of moral agency is actually and in fact considered a moral agent, however, depends on social context. For example, if an infant were stranded on an island entirely devoid of human life, and somehow survived to adulthood, that person would probably not be considered a moral agent. Morality is contextual—as long as there is no socialization pushing potential moral agents to become actual moral agents, it is possible for something to be capable of bearing moral responsibility without actually being a moral agent.

So, the moral capabilities of institutions are necessary, but not sufficient, conditions for their actually being considered moral agents. Some social expectation of moral behavior is required before institutions may be thought to bear moral responsibility. Social expectation answers the "how" question of institutional moral agency. Institutions are artifacts, and there is a non-trivial amount of self-determination that is necessary for an institution's existence. If society demands that institutions be held morally responsible for their actions, the institution can be created as a moral agent. The decision-making processes and organizational structure, the goals, and the principles of the institution can all be shaped with moral responsibility in mind. Institutions are set up to act, and society can legitimately and efficaciously demand that they act morally.

So far, institutions are potential moral agents, and it is up to people to demand that they bear moral responsibility for their actions. But why should society require actual moral agency of institutions? The simple answer is that the actions of institutions frequently have moral ramifications. Institutions act in ways that shape the lives of the individuals who are unquestionably a part of the moral realm. Because institutional action has moral import, we ought to expect moral responsibility from the agents.

There are, however, normative arguments that we ought not hold institutions morally responsible. One such argument comes from asking questions about the moral responsibility of individuals within institutions. Some philosophers argue that collective responsibility for moral action may allow for unfair moral burdens on individual members. I argue that accepting institutions as moral agents does not result in inappropriate responsibility judgments, and in fact allows for better and clearer ascriptions of moral responsibility. Although accepting institutional moral responsibility does not say much about collective responsibility in general, with an understanding that the moral agency assigned to an institution is not prone to individualists' concerns, more precise assessments of the moral responsibilities of agents in complicated institutional situations is possible.

First, a little more about the individualist's worries. In his seminal 1948 article on collective responsibility, H.D. Lewis argues that embracing collective moral responsibility would undermine the fundamental ethical assumption that no individual can be responsible for the actions of another—that "responsibility belongs essentially to the individual."[xxxi] The essence of Lewis' vehement rejection of collective responsibility is that acknowledging the moral agency of groups will nullify individual moral responsibility, and individuals will cease to be held to account for any action performed under the auspices of a group. As such, embracing collective responsibility is "not morality at all, but a repudiation of it."[xxxii]

More contemporary writers still subscribe to Lewis' argument. Jan Narveson argues that "to hold up a shadowy 'collective' as being the true responsible agent is to deflect responsibility from the only entities that can genuinely have it" —those responsibility-bearing entities being individual persons.[xxxiii] Similarly, Steven Sverdlik points to the Nazi massacre at Lidice, in which members of the Nazi party killed every male citizen of the village in retaliation for the murder of a German SS officer by a Czech partisan. Sverdlik explains that the massacre was wrong because the murdered villagers "were being held responsible for something they didn't do."[xxxiv]

All of these arguments for moral individualism tap into the same worries: Does allowing collective responsibility also allow for improper ascriptions of individual responsibility? Will

individual group members be beyond moral blame or praise? Might group members be inappropriately held responsible for the actions of the group? A "yes" answer to any of these questions would do irreparable damage to the idea that groups ought to be considered moral agents.

These objections are all directed at collective responsibility, understood with an unrestricted notion of what constitutes a collective. Narrowing the scope to focus on institutions, though, shows that at least some understandings of collective responsibility need not lead to any improper assumptions about the moral responsibility of individual group members. *Institutional* moral agency does not preclude holding individuals responsible for their actions, nor does it compel the release of individuals from their moral obligations.

The reasoning behind the above claim is one of logical nondistributability: just because an institution has been recognized as a moral agent does not mean that anything at all can be inferred about the individuals who make up the institution. The actions of individuals within institutions must be evaluated separately from the actions of the institution.

Peter French explains nondistributability with predicate logic: "If 'c' is a conglomerate and P is a responsibility predicate, moral or nonmoral, and if 'm' is a member of 'c,' then although Pc is justified, Pm may not be justified and can never be justified solel on the basis of Pc."[xxxv] So, a conglomerate may commit some action for which it is morally responsible. It may be that no individual member of the conglomerate is morally responsible for that action. And if there is an individual who is morally responsible for the action, her responsibility cannot be determined by the moral responsibility of the conglomerate alone.

Virginia Held also writes that it is good moral practice to assign responsibility to certain groups, even though "from such judgments…little follows about the responsibility of individual members of such groups; much more needs to be ascertained about which officials or executives are responsible for what before we can consider individual members of nations or corporations responsible."[xxxvi] It does not follow that individual members of a morally responsible collectivity are morally responsible for the actions of the collectivity.

It is possible to separate the moral responsibility of an institution from the moral responsibility of its component individuals because of the way an institution works. First, institutions are compatible with varying membership. Individuals within the institution can come and go with great frequency, without changing the institution's overall identity. The institution is not reducible to the sum of its individual members, so it is difficult to understand how moral claims about an institution could be reducible to moral claims about individuals.

Second, the organizational structure of an institution is designed to delineate just what kinds of individuals within the institution are responsible for what parts of the institution's activities, again representing a structural rift between the individual and the group. And the decision-making processes internal to any institution further serve to distinguish the individual from the group. By design, the organizational structures and internal decision-making processes of institutions separate the institution as a whole from its individual members. This separation allows for institutional action and means that the act of holding an institution morally responsible does not have any logical bearing on the moral responsibility of individual members.

In practice, these distinctions may not be black and white. It is often quite difficult to disentangle chains of moral action, particularly in hierarchical, personality-driven institutions. Many large corporations appear to be wholly dominated by the CEOs and chairpeople in charge; the Enron scandal was almost entirely focused on Ken Lay and Jeff Skilling. But there are good reasons for trying to determine which actions are institutional and which are individual. A theory of institutional moral agency that entails a logical separation between the institution and the individual members makes it possible to understand a kind of collective responsibility that does not fall victim to individualist's worries. Accepting that institutions and individual members of institutions may be discrete moral agents is a good first step to sorting out the tangles of responsibility that may lead to undesirable outcomes.

But how can one assess the moral responsibility of individuals within an institutional framework? Again, the theory of the nondistributability of moral responsibility between the institution and the individual seems imperiled by questions of moral interdependence: Are

individuals within an institution really independent moral agents, or is their volition impeded by their participation in a larger group? To what extent is individual independence constrained by the institution? And can morally exemplary action be expected from individuals within an institutional framework? I will argue that, like the question of praise and blame, the problem of moral independence can be clarified, if not entirely solved, by better understanding the characteristics of an institution that allow for moral nondistribution in the first place.

The moral agency of individuals within an institution can be assessed using the same criteria for moral responsibility described in the first chapter. However, the circumstances of institutional involvement will have some bearing on how each of these criteria is fulfilled, and especially on the degree to which each is fulfilled.

For an individual within an institution, causality is still the first requirement for moral responsibility. In order for an agent to be held morally responsible for some outcome, the action must be causally connected to that outcome, and must also be morally faulty. Take, for example, a worker at a factory that is (as an institution) knowingly producing unsafe cars, an action that (for the sake of this example) is unquestionably morally blameworthy. The particular worker is responsible for welding one small part to another small part—she is causally responsible for the unification of those two parts. She is also causally responsible, to a small degree, for the production of the car as a whole. In this case, one action is causally connected to two distinct outcomes: the unification of two car parts, and the creation of the car as a whole. The first outcome may or may not be morally blameworthy—perhaps it is this particular junction of parts that renders the entire car unsafe. The second outcome is morally blameworthy—the institution is knowingly producing unsafe cars. Clearly the question of individual moral agency cannot be understood on the basis of causality alone; in order to determine whether or not the factory worker is morally blameworthy, more needs to be known about her awareness, intent, and volition.

An individual may be aware of her individual actions without necessarily being aware of the actions of the institution as a whole. The autoworker is (presumably) aware that she is

welding parts together, and that her actions are contributing to the creation of vehicles that will be sold to the public. She may or may not know that the cars are unsafe. Even though the institution as a whole recognizes the safety concerns, that awareness may not be something factory-level workers are privy to. She also may or may not be in a position to question the safety of the cars. If she is employed by an established and well-reputed car manufacturer, she probably has little reason to question the institutional decision to produce the vehicle. Perhaps if the car manufacturer is well known for producing unsafe vehicles, or if she was instructed to use welding techniques that she knew to be shoddy or subpar, she would have reason to try and gain awareness of the institutional decisions. If she is acting with an awareness of the contribution she is making to the institutionally blameworthy action, her actions are also morally blameworthy.

In this example, awareness essentially covers intent as well. Once the welder has an awareness of the moral ramifications of her actions (i.e., that by doing her job and welding two auto parts together, she is contributing to the institutionally morally fautly action), it is hard to see how she could be unintentionally contributing to the production of a dangerous vehicle. Not all cases of moral agency are such that awareness entails intent, though, and it is important to establish to what extent the individual intended to participate in the actions of the institution.

Volition is where things begin to get interesting for the problem of individual moral independence within an institutional framework. Institutions can exact a terrific pull on individuals; members within an institution often feel as if they *must* act in certain ways. Acting otherwise means risking a job, a position, or membership. The autoworker has every reason to fear retribution if she speaks out against the decision of the institution. Some institutions are even designed to limit the volition of their participants. For example, many factories in urban China hire young girls from the country because, due to complicated Chinese residency laws, country residents are not permitted their own housing or health care in the cities. These laws allow the factories to keep their employees in dormitories and act as the sole providers of food and health care. If the employees do not toe the institutional line, they can find themselves without a salary, home, food, or healthcare. The level of volition in an institutional setting seems restricted, in

some cases to the point that there is hardly any volition at all. This is not to say that individuals within institutions are acting involuntarily. As defined in the first chapter, voluntary action is when the agent truly has no choice. To choose between acting according to the will of the institution and losing a job or being sent to jail is certainly not an easy choice, and I have no desire to downplay the severity of the restrictions on volition. But the fact remains that there is still choice, and thus individual action within an institutional context is not necessarily involuntary.

Another normative argument *against* understanding institutions as morally responsible agents comes not from individualism, but from worries about assignation of praise or blame. How can a collective be the subject of praise or blame? Is praise or blame of a collective equivalent to praising or blaming individual members? If blame cannot be placed squarely upon an collective, the individualist's worry about misplaced responsibility ascriptions resurfaces, and collective moral responsibility must be done away with.

Again, although institutional moral responsibility might not say much about collective responsibility as a whole, an account of institutional moral agency is helpful in solving this puzzle. Because institutions are singular actors, they can be the bearers of praise or blame, and there are good reasons for wanting to make moral judgments. And, although it may be difficult to maintain the absolute distinction between the individual and the institution in terms of being praised or blamed, the idea of logical nondistributability should be a first step toward making sure that moral responsibility is attributed appropriately and fairly.

Institutions, as singular entities, are capable of being subjects. Institutions can be commended by the press, rewarded by the government, even attacked by protesters. When the Center for Consumer Freedom takes out a full-page ad in the New York Times claiming "PETA kills animals,"[xxxvii] they are blaming PETA for perceived immoral action. In this case the immoral action is not actually killing animals; the Center for Consumer Freedom is actually blaming PETA for lying to supporters about its true institutional goals and principles.[xxxviii] This example,

taken not from the annals of history, but from an easily-forgotten ad in the morning paper, shows how ubiquitous the practice of praising or blaming institutions actually is.

It may be argued that praising or blaming an institution is a futile effort—can anything really be gained from judging the moral responsibilities of non-corporeal, non-human, non-feeling institutions? As Baron Edward Thurlow is rumored to have said about the corporations of the industrial revolution, "they have no soul to save and no body to incarcerate." So what is the point of allowing institutions into the community of moral agents?

An institution, as a singular entity, can accommodate not only varying membership, but varying levels of support or nonsupport from nonmembers, and assigning moral responsibility to institutions allows other moral agents to make decisions about their willingness to offer support. Virginia Held explains that making moral judgments about institutions "enable[s] us to reward acceptable corporations with our investments, or to shun travel to states that violate international norms."[xxxix] Institutions can be deprived of material support by investors, customers, or other nonmember participants. Parents may decide not to enroll their son in the Boy Scouts because they do not support the group's discriminatory practices—a moral judgment about an institution in this case leads to a gesture of nonsupport directed at the institution.

Furthermore, institutions are capable of change, and praise or blame of an institution may play a critical role in pushing institutional change. Institutions may be lacking bodies and souls, but, as I argued earlier, they do have goals and basic principles, both of which can be revised in response to praise or blame. For example, in the wake of a recent string of eating disorder–related deaths in Brazil, several fashion houses and modeling agencies have implemented minimum weight requirements for their models. The effort is, on one hand, an attempt to protect the health of their employees. But it is also a response to loud criticisms that skinny models promote unattainable aesthetic ideals—moral judgments asserting that the fashion industry is morally blameworthy for the spate of starvations.

Furthermore, given that institutions tend to have institutional histories, they have incentives to change that go beyond immediate concerns about praise or blame. Colby College,

for example, is quick to point out its relatively early decisions (compared to similar colleges) to admit female and black students. Having a commendable moral history is important to many institutions, and moral histories can be a source of institutional shame or pride. An institution may lack soul and body, but the presence of goals, principles, and history makes moral assessments of institutions worthwhile.

A final concern remains: for most institutions, institutional praise or blame translates directly into benefits or losses for individuals. Bad publicity for a business, for example, often means a fall in stock prices and the possibility of layoffs. In theory, the organizational structures and decision-making processes of institutions make the group entirely separate from its individual members. In practice, the distinction is not so simple.

In response to this concern, I say the practical problem of distinguishing the individual from the institution is not a problem for the *theory* of institutional moral agency. In fact, the practical difficulty of assigning responsibility is one reason to *embrace* a theory of institutional moral responsibility. If it is possible to accept that, in theory, institutions are moral agents distinct from their individual members, then it should also be possible to come up with systems of praise or blame reflecting that distinction. Thinking in terms of institutional moral responsibility can lead to a better way of praising or blaming individual group members.

It seems prudent to discuss scapegoating here, because sacrificing the individual for the sake of the institution is one practice that could benefit from an acceptance of institutional responsibility. Take as an example the Abu Ghraib prison scandal, the news of which broke in 2004. When photos surfaced depicting American soldiers torturing Iraqi prisoners, retribution was swift—but only for the soldiers in the photos. Even when it became clear that military top brass had been approving the use of torture in Iraq, there was no effort to affect change on an institutional level. On an individual level, it is appropriate that the soldiers in the photos be held morally responsible for their acts of torture. But by focusing entirely on the individual agents, the army failed to achieve any introspective institutional moral assessment. So here we see why a theory of institutional moral agency can actually help prevent problems like scapegoating.

When it comes to problems like employee compensation, or layoffs due to falling stock prices, the options for holding an institution responsible while sparing the individual members are more complicated. There does not seem to be any obvious way to blame or praise a business without rewarding or punishing individual employees. Fortunately, people tend to understand, when signing on to work for a company, that they are establishing a connection with an institution, and that their fortunes are linked to the business. This is particularly evident with major corporations who entice employees with stock options packages—in addition to being compensated for their individual performance, employees with stock options stand to benefit (or not) from the overall performance of the company. So even if, particularly in the case of businesses, it is nearly impossible to prevent individual ramifications of institutional moral agency judgments, it is at least possible to be aware of the linkages between the individual and the institution.

Because the connections are apparent, they are also evaluable and manipulable. Many businesses, particularly large corporations, hire in-house ethics officers or contract with outside ethical consulting firms. Compensation structures could easily be reviewed under the heading of ethics. If the company as a whole was to take a hit, and if individuals cannot be entirely separated from the institution, at least mechanisms can be instituted to make sure the individual outcome is fair. One way to do this might be to make sure pay is structured to reflect organizational and decision-making structures. In practice, it may not always be possible to maintain separation of individual and institutional moral agency. But the theoretical framework of moral nondistributability can at least be helpful in guiding the practically necessary connections.

Here I have shown that institutions, thanks to certain definitional characteristics, are capable of bearing moral responsibility. In order for institutions to be considered actual moral agents, people must consider them moral actors; like individuals, institutions must be socialized into the moral community. There are good reasons for wanting to hold institutions morally

responsible for their actions. Enough good reasons, in fact, that they outweigh normative worries about the potentially pernicious consequences of allowing institutional moral responsibility.

[xxii] Virginia Held, "Can a Random Collection of Individuals be Morally Responsible?" *Journal of Philosophy* 67 (July 1970), 471.

[xxiii] Peter A. French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984), 41.

[xxiv] ibid., 41

[xxv] ibid., 13

[xxvi] ibid., 13

[xxvii] ibid., 32

[xxviii] For more, see Hugh Curtler, ed., *Shame, Responsibility, and the Corporation* (New York: Haven Publications, 1986).

[xxix] For a very detailed description of the organization and decision-making structures of the Cold-War era Soviet Union, the United States, Britain, France, China, and NATO, see Scilla Elworthy and John Beyer, eds., *How Nuclear Weapons Decisions Are Made* (New York: St. Martin's Press, 1986).

[xxx] *Enron: The Smartest Guys in the Room*, DVD, directed by Alex Gibney, 2005.

[xxxi] H. D. Lewis, "Collective Responsibility," *Philosophy: The Journal of the Royal Institute of Philosophy* 23 (January, 1948), 3.

[xxxii] ibid., 7

[xxxiii] Jan Narveson, "Collective Responsibility," *Journal of Ethics: An International Philosophical Review* 6, no. 2 (2002), 184.

[xxxiv] Steven Sverdlik, "Collective Responsibility," *Philosophical Studies* 51 (1987), 64.

[xxxv] French, *Collective and Corporate Responsibility*, 15

[xxxvi] Virginia Held, "Group Responsibility for Ethnic Conflict," *Journal of Ethics: An International Philosophical Review* 6, no. 2 (2002), 157.

[xxxvii] *The New York Times,* 2007 (accessed 22 January).

[xxxviii] "PETA Kills Animals | PetaKillsAnimals.Com," http://www.petakillsanimals.com/

[xxxix] Held, *Group Responsibility for Ethnic Conflict*, 162.

Chapter 3: Situational Collections

In the previous chapter, I discussed institutions—complex, organized groups designed to make decisions and achieve goals—and how they and their members can be understood as morally responsible agents. But institutions are an exceptional kind of group, given their unique structure and design. Less organized groups, groups lacking any internal decision-making process, may also be candidates for collective moral responsibility. If a group of people on a subway could work together to stop a violent crime that any one individual would be incapable of stopping, that group may well be responsible for its failure to act *as a group*.[xl] It is possible to imagine many actions that require the input of more than one individual agent for success, and it is easy to see that this type of action could have moral relevance. If an action requiring many individuals for completion has a praise- or blameworthy outcome, then there must be some way of understanding the moral responsibility of that group and the individuals within it.

I will argue that, even using the best possible description of these kinds of groups, it is unclear whether or not they can satisfy the minimum requirements for moral agency. Then, I will show that, charitably assuming these groups can act as singular moral agents, an understanding of collective moral agency will result in either a morally unacceptable or a philosophically trivial account of moral responsibility. From there, I will posit a description of the groups in question based on cooperative, rather than collective action, and suggest that individual actors have a share in the responsibility for the cooperative action.

The kinds of groups in question are collections like the group of people on a subway, who must act together to stop a violent crime. Virginia Held names these groups "random collection[s] of individuals."[xli] For Held, a random collection is "distinguishable by some characteristics from the set of all persons, but lacking a decision-method for taking action."[xlii] Held does not lay out any criteria by which collections may be distinguished; rather, her collections are defined by their absent decision-making method and "some characteristics" setting the group of individuals

apart.[xliii] Held gives the example of "passengers on a train or pedestrians on a sidewalk,"[xliv] who are distinguishable based on their shared location or a similar destination. There is nothing bringing these people together except for coincidence of setting, but that coincidence is enough to distinguish that group from the set of all people.

Held's definition must be stronger to be useful—according to her description a group made up of my grandmother, the rapper 50-Cent, and Napolean Bonaparte is a random collection, simply because I have distinguished those individuals from the set of all persons (their distinguishing feature is that I happen to associate them, as opposed to anyone else in the world, with the word "random") and because they have no decision-method for action. Perhaps more importantly, if any criterion can be used to separate a random collection from the universe of every individual, many random collections could emerge that look nothing at all like the groups in Held's examples. The collection of blue-eyed people is a set of people who are distinguished from the set of all people and who do not possess a decision-making method, but that group differs so significantly from Held's bunch of passengers in a train-car that it seems illogical to attempt a moral understanding of collective responsibility that encompasses them both.

I think the strongest account of these kinds of collections will emphasize the shared situation of the group. Focusing on shared situations spares the group from true randomness, while maintaining Held's sense that the group is constituted because of the accident of coincidence. "Situation" can have many meanings; the simplest is "the position of something in relation to other things" or "position of a person with regard to circumstances."[xlv] However, in this sense, "situation" is a technical term. This is the situation of live newscasts or police reports, "The Situation" as opposed to "situate" or "situated," and "The Situation" carries with it a connotation of crisis and resolution. A situation is something troublesome, something that calls for action or amelioration. In this sense, all of Held's random collections are people who are distinguished based on a shared situation, and the moral evaluation of these collections will be based on their response to the situation at hand. Instead of calling these groups "random collections," then, they may be more efficiently thought of as situational collections.

So, to sharpen Held's definition, a situational collection is "a set of persons distinguishable by *their shared situation* from the set of all persons, but lacking a decision method for taking action that is distinguishable from such decision methods, if there are any, as are possessed by all persons."[xlvi] Also, sitcols are what Peter French calls "aggregate," not "conglomerate."[xlvii] Unlike institutions, whose identity persists through changes in membership, a sitcol changes when individual members enter or exit the scene. French writes, "[a] change in an aggregate's membership will always entail a change in the identity of the collection,"[xlviii] and this will always be true of a sitcol. If someone leaves the group, the situation has changed. If someone joins the group, the situation has changed. Because membership is part of the group's situation, and because the sitcol is defined by its situation, a change in membership entails a change in identity, and thus the sitcol can be understood as an aggregate. The aggregate nature of situational collectives will have some bearing on the responsibility of those collectives and on the people in them.

So, sitcols are aggregate collectives in which a group of individuals can be distinguished from the set of all individuals on the basis of a shared situation. They possess no internal decision-making process, and their identity changes with membership change. The examples from Held's "Can a Random Collection of Individuals Be Morally Responsible?" are all examples of situational collectives[xlix], as are French's examples of aggregate collectivities.[l] But even given the most precise definition of situational collections, it is not clear that sitcols, as collectives, are capable of bearing moral responsibility.

Held says that these collections can be morally responsible—sometimes: "When the action called for in a given situation is obvious to the reasonable person, and when the expected outcome of the action is clearly favorable, a random collection of individuals may be held responsible for not taking a collective action."[li] So, as long as the sitcol is in a situation that calls for action, as long as the action called for is obvious, and as long as the outcome would be favorable, sitcols can be held responsible. In addition, "a random collection can be held morally

responsible for failing to make a decision on which action to take—for failing, that is, to adopt a decision method."[lii] Essentially, a sitcol can be held responsible for failing to act as a group.

Peter French, on the other hand, argues that a sitcol can *never* be the bearer of moral responsibility, because a sitcol is an aggregate and "moral responsibility predicates cannot be legitimately ascribed to aggregate collectivities."[liii] He explains that, even when talking about actions that could not have an individual subject—he uses the example of a mob rioting, because a mob can riot, but mob member x cannot riot—the moral responsibility will always be predicated to individual actors rather than collectivities.[liv]

What about the four criteria for moral responsibility developed in the first chapter? Does Held's assertion that these groups can (sometimes) bear moral responsibility hold up to the requirements for moral agency? Maybe French is wrong and (aggregate) collective actions *can* have collective predicates.

Causality first: there are certain actions, like moving a piano up a flight of stairs, which truly cannot be performed by a single person. With these kinds of actions, which cannot be completed without the participation of more than one person, it may well be helpful to collapse the many individual causers into one collective causer. If a group of motorists, all stuck on the same stretch of road by a large fallen tree, work together to move the tree and clear the road, it may be possible to consider the group, as a collective, causally responsible for the moving of the tree. These cases show why one might want collective action, and what it might look like. But it is not always easy to explain *how* these collective actions occur. Unlike institutions, which are explicitly designed to be able to act institutionally, situation-based collections are not made to act. Any account of collective action would have to explain *how* a collective could be considered to act collectively, and such an explanation does not seem readily available. Sitcols *might* be capable of collective action, but more would need to be explained to understand how.

The question of collective awareness for sitcols is also is answered with a "sort of." Held writes that collective action does not imply "the existence of an inexplicable group awareness over and above the awareness of its individual members, only that we are sometimes entitled to

say 'Random collection *R* is aware that *p*'." [lv] She gives an example of a group of people, all of whom are both "normal" and are standing (as opposed to sitting). Held says it may be true that the group of people are aware that they are standing, even without knowing that any individual member of that group had such an awareness.[lvi] Held does not want to posit some supermind, capable of moral awareness and existing singularly and apart from the minds of the individual group members. Rather, she argues that, in some situations, it is plausible to make true statements about a group of people being aware, without necessarily being able to make assertions about which individual group member is aware.

So far, it looks like sitcols might be able to act collectively—sometimes. And sitcols might be collectively aware—sometimes. Assuming that "sometimes" is enough to satisfy the responsibility criteria, what about intention? There is plenty of literature on "we-intentions,"[lvii] and shared intentions,[lviii] although it is not always clear if this sharing of intentions is the same as collectively holding an intention. A group of individuals sharing in one intention is not necessarily the same as a collective, as a single agent, intending. Rather than each person literally sharing the same intention (as one might imagine every individual in a group sharing the same coffeecake), Velleman seeks to explain how many individual intentions can "add up to a single token of intention, jointly held."[lix] Although shared intentions may well be useful for explaining how a situational collection can jointly intend, it is not clear that embracing shared intentions will entail accepting collective intention for sitcols.

On the other side of the sitcol-intention argument is Peter French, who says that aggregate collectivities are not intentional, without delving much into why they are not (his focus is more on conglomerate collectivities).[lx] And Christopher Kutz ponders the importance of participatory intentions, which are not collective intentions held by a group, but singular intentions to bring about a collective outcome. So, in the case of a sitcol working together to remove a fallen tree from a road, each individual has a participatory intention. *Each* intends to participate in moving the tree, thus each intends to *individually* contribute to the accomplishment of some collective outcome. For Kutz it is participatory intentions, rather than a single collective

intention, that are important for explaining collective action.[lxi] Even when actions are taken collectively, there is no collective intent, and thus no collective responsibility.

Unlike intention, there is no large body of writing on collective volition that might be relevant to determining the ability of situational collections to meet the last criterion for moral responsibility. Can groups act voluntarily?, though, seems likely to have the same kind of equivocal answer as questions about causality, awareness, and intention: "maybe," sometimes," or "yes and no." Overall, then, it is not clear, based on the lines of argument explored above, that sitcols conclusively can or cannot bear moral responsibility. Assuming they can, however, it is also unclear why it would ever be useful to assign moral responsibility to a situational collection.

If responsibility is assigned to a situational collective, it will distribute equally to its members. Held suggests that the moral responsibility of such a collective "seems to be distributive," writing: "If a random collection *R* can be represented as a set equivalent, say, to *M* & *N* & *Q*, then, if *R* is morally responsible, we would seem to be able to conclude that *M* is morally responsible & *N* is morally responsible & *Q* is morally responsible."[lxii] The distribution of moral responsibility is metaphysically murky. It is unclear whether the individuals each get all of the responsibility, if they each get a piece of it, or perhaps "partake" of it. Questions reminiscent of Aristotle's critique of Plato's forms seem to come up in response to the notion of responsibility distribution. Distribution does not seem like a specific enough idea to be widely accepted.

Furthermore, if responsibility distributes evenly, situational collective responsibility is trivial. Once distributability has been accepted, collective responsibility does no theoretical work. Held herself says, "in saying that the moral responsibility of a random collection is distributive, we may not be saying very much."[lxiii] Because the moral responsibility is parceled out evenly among members of situational collectives, there seems to be no reason to attempt to save the possibility of separate moral agency for the collective itself. Even if it were affirmatively

plausible to assign collective responsibility to situational collections, it does not seem practically useful.

The above argument does not do away with the intuitive desire to hold situational collections collectively responsible for their actions. Virginia Held describes several situations that appear to warrant a theory of situational collective responsibility. Her examples include a group of people on a train who could work together to stop a deadly beating, and three people who spend too much time arguing about a course of action to save the life of a person trapped beneath heavy beams.[lxiv] Also, in cases involving mobs or rioting, it may be especially tempting to view the group as the responsible agent. The idea of "mob mentality" suggests that people act differently, and are perhaps not entirely in control of their actions, when involved with crowds or groups. Because mobs seem to be able to change the individuals within them, it may make sense to want to evaluate such groups independently of their individual members.

In the examples above, the temptation to place moral responsibility squarely on the collective stems from discomfort with holding individuals responsible. Particularly in the cases of mobs or riots, it is an unease with blaming individual mob members that leads to a temptation to blame the group as a whole. But this is just the kind of pernicious blaming that leads to harsh criticisms of collective responsibility. Assigning moral responsibility to ad-hoc "superagents" in order to exculpate individuals is exactly the kind of "barbarous" non-ethics H.D. Lewis was worrying about in the 1940s.[lxv]

Assigning responsibility to sitcols is different than assigning responsibility to institutions, which is why situational collective responsibility, but not institutional collective responsibility, falls victim to arguments like Lewis'. Institutions are stable. They may be artifacts, but they are long-lasting artifacts designed to perform the kinds of actions that warrant questions of moral responsibility. And, as I pointed out earlier, institutions have goals and histories and the ability to change. As artifacts, their stability and design make them fair game for assignations of moral responsibility. Sitcols, on the other hand, are fleeting and mercurial. To pin responsibility on a sitcol is to hold responsible a particular group that existed at a particular time and place and

which is probably nonexistent (in collective form) at the time of the blaming. Sitcols are not collectives that persist, and as such, are not collectives that can react to things like praise or blame. This instability makes sitcols markedly different than institutions, and is the reason situational collective responsibility is problematic.

So, even if one *can* hold situational collections morally responsible (although I'm not sure one can), and even if holding situational collections morally responsible would have some nontrivial implications for understanding morally responsibility more generally (although I'm not sure it does), the ethical implications of accepting situational collectives as potential moral agents seem unacceptable. What, then, to do about all the examples seeming to warrant an account of situational collective responsibility?

So far, I have posed a challenge: how to respond to examples of non-institutional collections, united by a shared situation, who may appear to behave collectively, or fail to constitute themselves as collective? I have suggested that collective action and collective responsibility are not adequate for mapping out moral responsibility ascriptions for these kinds of collectives. I will now attempt to show that two theories, Michael Bratman's shared cooperative activity,[lxvi] and Larry May's shared responsibility,[lxvii] are better models than collective action and collective responsibility in these cases.

Thinking in terms of cooperation preserves the desire to think about actions that could not have been accomplished without the participation of many individuals, without the inefficacious or morally pernicious assignation of collective moral responsibility. Shared cooperative action refers to a group of individuals, each working as individual agents to produce some outcome requiring many actors for its achievement. Michael Bratman says that "shared cooperative activities" are characterized by multiple individual actors who exhibit "mutual responsiveness," "commitment to the joint activity," and "commitment to mutual support."[lxviii] He argues that, in cooperative activity, all actors are acting and responding to the intentions and actions of others. All actors intend that the act be completed, and be completed by cooperative action, and each

actor is committed to supporting each other actor in the pursuit of completing the cooperative action.[lxix]

Applying Bratman's cooperative action model to situational collections is helpful in explaining the responsibility of the individuals within those situations without appealing to collective moral agency. Cooperative, rather than collective, action allows for different ascriptions of responsibility to members of a situational collection. In any given situation, a number of actions may be taking place. Certain members of the group may work together and act cooperatively, but it is not necessary that each member of a situational collection participate in cooperative action. Individuals may dissent, may decide to do things differently, or may try to escape the situation entirely. A group of people at a bus stop, confronted with a crisis requiring action for its moral resolution (for example, a bank robbery across the street) may choose to act in total cooperation: every person at the bus stop may agree to help stop the robbery, and they may quickly devise a cooperative plan. Or, one persona may argue that it is not the business of anyone outside of the bank, the robbers, and the police, and refuse to take part in any cooperative vigilantism. Another person might decide that this particular bus stop is overly problematic and hail a cab. The specifics here are less important than the point that, in any given situational collection, myriad activities—both individual and cooperative—may be taking place. The true task in evaluating moral responsibility in situational collection contexts is to decide which of this multiplicity of actions contributed to the outcome in question, and which individual agent or agents performed or cooperated in that action in an aware, intentional, and voluntary way.

Shared cooperative action, then, accounts for the intuition that, in certain situations, a group must act together to achieve a morally desirable outcome. It also preserves individual moral responsibility by allowing for individual action within the situational collection. Cooperative action accounts for the first requirement of moral responsibility: causality. Individuals can share in causality as part of a cooperative action; for example, many individuals may cooperate to stop an assault on a subway. They can be causally responsible for failing to participate in a morally good cooperative action; a person who sits out and fails to help is

responsible for her own nonparticipation. Collective action does not account for these (not always straightforward) examples of causality in situational collections; shared cooperative activity does.

As for the other criteria for moral responsibility, once the causality tangle has been sorted through, individual awareness, intention, and volition apply to members of a situational collective in the same way as they would apply to individuals. There is no collective to which these requirements may be assigned. As was the case with individuals who are members of institutions, however, it is important to think about the ways membership in a situational collection may effect individual responsibility. This is particularly true in the case of volition; there may be many reasons for an individual in a situational collection to feel as though their volition is restricted. Kutz's analysis of participatory intentions helpfully distinguish between two parts of collective activity: the individual action and the collective outcome.[lxx] Since volition concerns the ability of an agent to choose a course of action or have control over an outcome, Kutz's distinction is a useful starting point for thinking about volition. First, an individual engaging cooperatively in an activity is performing some individual action—say, lifting one end of a large, heavy object. The individual has control over that action in the same way she has control over any other personal activity. Given a particular situation, an individual chooses some action; the fact that other people are involved is a part of the situation the way the available sunlight, or the presence of furniture, are. The situation helps guide the course of action, but the individual choice to act is voluntary.

There are, however, many reasons for an individual in a situational collection to feel as though their volition is restricted, or has been restricted. One may be dealing with coercion or peer-pressure, or feel outnumbered or threatened. Or, one may be overwhelmed by feelings of solidarity or mob mentality. Going back to Aristotle, though, coercion does not lead to involuntary action by the coerced. It may call for a lesser degree of volition, and a correspondingly lower degree of responsibility, but peer pressure cases do not appear to result in impracticality.

The second part of Kutz's description of collective action is the collective outcome (as opposed to the individual action). The collective outcome is frequently something about which an

agent can truthfully say "I had no control." In the case of truly collective outcomes, things which could only have happened with the participation of more than one person, the outcome may well be outside the individual's control.

Shared cooperative activity works better than collective action as a model for understanding the responsibility of people involved in situational collections. Similarly, shared responsibility works better than collective responsibility for thinking about how to praise or blame members of a situational collection.

Larry May distinguishes three kinds of responsibility: total responsibility, no responsibility, and partial responsibility.[lxxi] In the case of institutions, the institution is assigned total responsibility for the institutional actions, and the individuals within the institution are assigned total responsibility for their own individual actions. Worries about collective responsibility for situational collectives revolved around the fear that individuals would bear no responsibility, and that the placing of total responsibility on the collection would not actually accomplish any moral good. So, what we need is a conception of shared responsibility. If individuals in a situation are engaging in shared cooperative action, they are sharing in the creation of a particular resolution. Thus, they must share the responsibility for that resolution.

How does this work? May asserts that the responsibility for a harm may be divided, paying attention to "the way in which each of us interacts with others," rather than focusing on the individual outside of situational context.[lxxii] This is why it was so important, in the first chapter, to put forth a set of criteria for moral responsibility that allowed for degrees of responsibility. In the case of a situational collection, each individual participant may be responsible for an outcome to some degree; the responsibility is shared.

It is possible to argue that "shared responsibility" is just a different way of looking at collective responsibility. The above account of responsibility as shared among individuals, each of whom are responsible to a degree for some action, could be described differently: total responsibility was assigned to the collection, and that responsibility was shared among the members of the collection. But that is not an appropriate formulation of the shared responsibility

model; it is backwards. Rather than taking "total responsibility" and divvying up, like a pie, shared responsibility deals with shares of responsibility. Each individual sharing in the responsibility contributes a degree of responsibility to the outcome as a whole: think of poker players anteing into a game. There is no "total pot" until all the antes have been collected, similarly there is no "total responsibility" until each individual's degree of responsibility has been accounted for. The true meaning of shared responsibility, then, is starkly different from collective responsibility.[lxxiii]

To see why this must be, recall that not every individual who is part of a situational collective must be involved in a cooperative action associated with the situation. Trying to suggest that an entire situational collection is sharing in some total collective responsibility, assigned to the group as a collective, would make it difficult to account for the outlying individuals who did not contribute cooperatively to the outcome in question. The familiar example of passengers on a subway illustrates why collective responsibility divvyed amongst individuals is not a workable model. If everyone in the subway cooperated to subdue the assailant and end the violence, then perhaps it would be fair to say that the group is collectively responsible (in the praiseworthy sense) and that the responsibility can be evenly divided amongst the individual members. However, say two able-bodied people sit out and watch the rest of the car stop the violence, and one young woman holding an infant offers to help, but ends up protecting her child instead . The situational collection would have to be arbitrarily divided: the collection of people who did cooperate to stop the violence the collection of people who did not cooperate, but could have; the collection of people (one person) who wanted to cooperate, but could not; and the collection of people (also one person) who lacked the fine motor skills to do anything other than bawl. From there, collective responsibility would be assigned to the collection of people who did not cooperate but could have (negative, blameworthy responsibility) and to the collection of people who did cooperate (positive, praiseworthy responsibility). Such a parceling out works, but since situational collections are separated out from the set of all people by virtue of some shared situation, this ad hoc division does not make sense.

Looking at the same example in terms of shared responsibility is clearer. Each individual who participated in stopping the assault is responsible, to a degree, for the stopping of that assault. Each individual who did not participate is not responsible for that same outcome, nor are the woman and her child. If, however, the cooperative action failed, and the assault went on, the responsibility is shared by different individuals within the collective. Those able-bodied individuals who failed to participate in the cooperative action would each be responsible, to a degree, for failing to cooperate. A strict adherent to the rules of responsibility would say that the woman with the infant is responsible, to a degree, albeit a small one given the restrictions on her volition (a choice between protecting your child and engaging in cooperative policework does not seem all that voluntary) and the good intentions evidenced by her offer of help.

The model of shared responsibility, then, explains how responsibility is assigned to members of situational collectives better than the model of collective responsibility. This is not, of course, to suggest that there is no collective responsibility: I spent the entire last chapter defending collective responsibility in the context of institutional action. Collective responsibility is just not the best way to explain how moral responsibility may be understood for, at least, situational collections.

There is another benefit, too, to thinking about shared responsibility. This account addresses a critical gap in my explanation of institutional responsibility. The institution receives full responsibility for its own action, and the individual receives full responsibility for her own action. Yet there is an important sense in which individuals within institutions cooperate to resolve situations. There are many actions that occur in an institutional setting that could not be accomplished without the joint participation of many individuals. An account of shared responsibility can complement the previous understanding of institutional moral responsibility to show how individuals can also have shares in the responsibility of actions taken as part of an institution (actions which may have had an effect on the institutional action, but which were not necessarily the same as the institutional action and which also cannot be construed as individual action).

There is no obvious way, then, to describe situational collections in a way that allows them to be singular, collective moral agents. It makes more sense, in these cases, to think about moral responsibility based on models of shared cooperative activity and shared responsibility. Unlike institutions, which could be described as singularly morally responsible, situational collections are not the kind of group that can accommodate talk of collective moral responsibility.

[xl] Example from Virginia Held, "Can a Random Collection of Individuals be Morally Responsible?" *Journal of Philosophy* 67 (July 1970), 471–480.

[xli] ibid.

[xlii] ibid.

[xliii] ibid.

[xliv] ibid.

[xlv] "Oxford English Dictionary: Situation," http://0-dictionary.oed.com.library.colby.edu/cgi/entry/50225868?single=1&query_type=word&queryword=situation&first=1&max_to_show=10 (accessed March 12, 2007).

[xlvi] Held, *Can a Random Collection of Individuals be Morally Responsible?* I have replaced "some characteristics" with "their shared situation."

[xlvii] Peter A. French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984), 13.

[xlviii] ibid.

[xlix] Held, *Can a Random Collection of Individuals be Morally Responsible?*, 471–480.

[l] French, *Collective and Corporate Responsibility,* 13.

[li] Held, *Can a Random Collection of Individuals be Morally Responsible?*, 475.

[lii] ibid.

[liii] French, *Collective and Corporate Responsibility*, 13.

[liv] ibid.

[lv] Held, *Can a Random Collection of Individuals be Morally Responsible?* 479.

[lvi] ibid.

[lvii] See, for example, Raimo Tuomela, "We Will Do It: An Analysis of Group Intentions," *Philosophy and Phenomenological Research* (June 1991), 249–277.

[lviii] Notably, Michael Bratman, "Shared Intention," *Ethics* 104, no. 1, (October 1993), 97–113; and J. David Velleman, "How to Share an Intention," *Philosophy and Phenomenological Research* 57, no. 1 (March 1997), 29–50.

[lix] Velleman, *How to Share an Intention,* 31.

[lx] French, *Collective and Corporate Responsibility.*

[lxi] Christopher Kutz, "Acting Together," *Philosophy and Phenomenological Research* 61, no. 1 (July 2000), 1–31, http://links.jstor.org/sici?sici=0031-8205%28200007%2961%3A1%3C1%3AAT%3E2.0.CO%3B2-%23.

[lxii] Held, *Can a Random Collection of Individuals be Morally Responsible?* 477.

[lxiii] ibid.

[lxiv] ibid.

[lxv] H. D. Lewis, "Collective Responsibility," *Philosophy: The Journal of the Royal Institute of Philosophy* 23 (January, 1948), 3–18.

[lxvi] Michael Bratman, "Shared Cooperative Activity," *Philosophical Review* 101, no. 2 (April 1992), 327–341.

[lxvii] Larry May, *Sharing Responsibility* (Chicago: University of Chicago Press, 1992), 204.

[lxviii] Bratman, *Shared Cooperative Activity*, 328.

[lxix] ibid., 328–329

[lxx] Kutz, *Acting Together*, 10.

[lxxi] May, *Sharing Responsibility*, 204.

[lxxii] ibid., 38.

[lxxiii] This account of shared responsibility does not exactly jive with May's description. However, I think this understanding of shared responsibility as cumulative, rather than distributive, is more useful and less prone to attack.

CHAPTER 4: SHARED ATTITUDE COMMUNITIES

So far, I have written about individual moral responsibility, the collective responsibility of institutions, and a non-collective model of responsibility for groups united by a shared situation. But there is an important kind of group that has been left out of the discussion. History is burdened by harms attributed to whole classes of people. Sometimes blame is placed on a nation, as in "The German people are to blame for the Holocaust," or "as an American, I fee guilty for the war in Iraq." Sometimes it is placed on an ethnic group: "the Rwandan Hutus are responsible for the genocide of Tutsis and moderate Hutus in 1994." Or, blame can be placed on a group of 'ists or 'phobes: "Racists are responsible for the legacy of oppression and violence against blacks in America," or "homophobes everywhere are to blame for the brutal death of Matthew Shepard." Given that events like the Holocaust and the United States' involvement in Vietnam catalyzed philosophical discussion of the possibility of collective moral responsibility, it is not possible to leave these sorts of blame statements unexamined in a paper about collective responsibility. My goal for this chapter is to evaluate these blame statements to explain their moral resonance and meaning. I will posit the existence of shared attitude communities, and endeavor to describe these groups in a way that clarifies the scope and degree of blame, both individual and collective, being attributed in blame statements like the ones above.

Looking first at the subject of the blame statements: who is really being blamed? The groups in question are special, and can be considered one kind of group, like an institution or a situational collection, because their individual members share a particular attitude. Although it is tempting to say that the subjects are national or ethnic groups, and although there are certainly many interesting discussions of national responsibility[lxxiv] and the responsibilities of ethnic groups,[lxxv] there is more to these statements than national or ethnic identity. Statements about racism are included in the relevant blame statements, for example, and racism certainly defies national or ethnic boundaries.

Larry May is helpful here; his account of racism seeks to hold people responsible for shared attitudes[lxxvi], and his model is good for making sense of the subjects of these statements. All of the statements of blame from the introduction can be recast to incorporate shared attitudes: "the German people, characterized by their attitudes of indifference, are to blame for the Holocaust," "the Rwandan Hutus, united by a prejudicial attitude toward Tutsis, are responsible for the genocide of 1994," and "racists, or, the group of people who share certain racist attitudes, are responsible for the legacy of oppression and violence against blacks in America." These rewrites are, if syntactically a bit clumsy, certainly not violating the sense or spirit of any of the blame statements. Recasting statements of blame to make explicit the underlying shared attitudes helps explain why "the German people," "Hutus," and "racists" seem to belong in the same category of group.

So what is an attitude? According to May, "Attitudes are not mere cognitive states, but they are also affective states in which a person is, under normal circumstances, moved to behave in various ways as a result of having a particular attitude."[lxxvii] Attitudes are states that spur action, and their maintenance is considered voluntary. They have concomitant sets of predictable action. If a person has a particular attitude, it is possible to predict that, in a certain situation, a person with a given attitude will act in a particular way. The having of an attitude entails a particular way of acting in the world. The subject of these blame statements, then, are groups whose members all have some attitude in common.

What is still unclear, then, is whether the shared attitude community as a whole is being said to bear the responsibility, or whether it is just individuals who are being blamed. So, to say that the subject of blame statements like "the German people are responsible for the Holocaust" is a shared attitude community—in this case, the community of people who shared an attitude of complacency with regard to the actions of the Third Reich—does not say very much about the moral content of the blame statement as a whole. More needs to be done to explain *who* exactly is being blamed, and for what. To this end, I argue that the blame statements in question are far more complex than their simple grammatical arrangement might suggest. Rather than assigning

blame to one person or one group for one harm, there are three levels of responsibility, three kinds of subjects, and three kinds of harms implicit in the statements in question. There are three ways to read each blame statement, and each reading contributes to the overall account of the responsibility of shared attitude communities. In this way, reading these blame statements is like looking at a 3-D image on the back of a cereal box through 3-D glasses. The image itself is composed of distinct lines, but with multi-colored lenses, you see the image as a whole. In order to understand the complexity of the blame statements in question, I will spend the rest of this chapter looking at them without the benefit of 3-D glasses, in order to see distinctly the lines making up the whole.

The Community Itself

Can the shared attitude community itself be considered a singular agent? Perhaps, like a situational collection, each community member is being blamed for her own actions, carried out cooperatively. In many ways, the shared attitude community looks like a situational collection. The shared attitudes uniting all the members of a shared attitude community could be considered a shared situation—it is a fact of my situation that I possess certain affective attitudes. Also like a situational collection, a shared attitude community lacks the decision-making procedures and organizational structures that allow for collective action. Using the situational collection model of responsibility, with its emphasis on cooperative action, does not seem likely, however. The barriers to cooperation for an entire shared attitude community are quite high; a group of people united by a common attitude does not seem like the kind of group that can meet the requirements for truly cooperative action.

The institutional model of collective responsibility is a potentially more likely means of explaining the kind of moral agency a shared attitude community may evince. Although a shared attitude community has no formal decision-making process or organizational structure, and although it is not the sort of group that can act as a singular agent, there are many ways in which a shared attitude community looks like a single entity. A shared attitude community, like an

institution, persists through membership change, and thus is a conglomerate, not an aggregate group, so it is more than the mere sum of its constituent members.[lxxviii] And a shared attitude community, again like an institution, has goals, a history, and basic underlying principles. The community does have some efficacy in that it can provide impetus for individual action. An individual within the community may feel more inclied to act on an attitude knowing that there is a community holding similar attitudes. The shared attitude community does not *act* in the sense that an institution acts, but it is not inert.

May has a helpful explanation of a shared attitude community as more than the sum of its parts. He writes: "The shared attitudes within a community come, over time, to create a shared identification, a shared feeling of solidarity."[lxxix] This explanation has intuitive appeal; even the most mercurial of shared attitude communities are bound by some feeling of solidarity. The evidence for this is mundane—a feeling of connectedness to a particularly well-written op-ed piece explaining an attitude you also hold, a willingness to identify yourself based on attitudes for social networking purposes, a desire to join clubs and organized institutions where attitudes may be developed or explored—all of these feelings suggest that there is something unique about shared attitude communities that makes them singular entities, albeit in a shifting and hazy sense.

Neither the situational collective nor the institutional models seem to exactly explain what kind of moral entity a shared attitude community is. A shared attitude community is not singular in the sense that it is an acting, causal agent like an institution. But the existence of a community may make it more likely that an individual will act in a materially harmful way and, given that the shared attitude community is not just describable based on the compiled description of its constituent parts, there must be some sense in which the community itself bears responsibility for harm.

The singular agency of the shared attitude community violates the four requirements for moral responsibility—it does not seem possible to describe a shared attitude community as causal, aware, intentional, or capable of volition. So it is not entirely proper to read the blame statements in question as being assignations of responsibility to a shared attitude community as a

singular agent. However, although it might not be an accurate description, it may be instrumentally useful to think of shared attitude communities as singular agents only insofar as it is not possible to completely understand blame statements as aggregated ascriptions of the responsibilities of individual community members. So, it would not be appropriate to hold a shared attitude community responsible in quite the same way as an institution: one could not blame an attitude community and expect it, as a singular agent, to respond and become a better moral actor. However, it is helpful for understanding the blame statements to realize that there is an important sense in which the responsibility of a shared attitude community cannot just be a totting up of individual moral responsibilities.

The Individual in the Absence of Causality

Another reading of the blame statements in question places the blame on the individual attitude holder *just* because they hold an attitude. So, to say "racists are responsible for the legacy of intolerance and oppression of blacks in America" is to blame individual racists for having racist attitudes. The blame here, like the blame being ascribed to the community itself, is not tied to any causality. According to this reading a person in a coma, incapable of movement and lacking the capacity to communicate, could be a morally blameworthy attitude-holder; one need not act upon the attitude in order to be morally responsible for the harm produced in the existence of the shared attitude community.

At first glance, the reading seems faulty—how can a person incapable of action or communication be contributing to harm? Larry May explains that the real harm at issue is the *possibility* of doing direct material damage. Certain attitudes, he says, are a kind of moral recklessness, and people who hold morally reckless attitudes should not be exculpated just because they were lucky enough to not act upon their attitudes in a directly harmful way. For May, racism is the morally reckless attitude exemplar, and he writes:

> Those who hold racist attitudes do not *do* anything that could be said to stand in the causal chain leading to racially motivated violence. But insofar as these people do not try

> to decrease the chances of such violence by changing their own attitudes, given that similar attitudes in others have produced harm, they demonstrate a kind of moral recklessness…which implicates them in racially motivated violence.[lxxx]

May's argument is that the only thing preventing people who hold harmful attitudes from *actually* causally contributing to physical or metrial harm is luck, and luck ought not be relevant for our moral decision-making: "The racist who does not cause harm is responsible because he or she shares in the attitudes and dispositions that, but for good luck, would cause harm."[lxxxi] Although May's argument is compelling, his moral recklessness account fails to satisfy the first of my conditions for moral responsibility. *Potentially* causing direct physical harm is not the same as causing direct physical harm. Unless the agent is actually causally linked to the physically or materially damaging actions, then the agent must be causing something other than the material harms associated with her attitudes.

There are other ways in which individuals may be considered blameworthy, if not causally responsible, for physical or material harm that emerges from the existence of a shared attitude community. Linda Radzik proposes the "Blood Money" theory of responsibility, in which she accounts for "benefiting from the harms one's group inflicts on others."[lxxxii] According to Radzik, an individual "will count as responsible in [the Blood Money sense] even if he himself has done nothing to create, encourage, or perpetuate the system of discrimination."[lxxxiii] For example, I do not consider myself a racist and I would not say I hold racist attitudes. But I certainly benefit from the legacy of racism in America—I recognize that there is plenty about my life that is easy for me not because of the color of my skin, but because of the persistence of a racist system of oppression that has not been successfully eradicated. My lifestyle is not without moral taint, because I benefit from the continuing existence of harmful attitudes. Radzik explains that responsibility, in the Blood Money case, is a case of "moral debt."[lxxxiv] Comparing moral debt to financial debt, she suggests that moral debt "can be incurred through no fault of one's own."[lxxxv] As a result, an individual who benefits from the harmful attitudes of others has "a duty to settle accounts, to return the 'stolen goods' that he has inadvertently received."[lxxxvi] Thus, even

if there is no causal link to any particular harm, there is some sense in which an individual ought to act as though she is responsible. She can take responsibility, even if she is not being held responsible.[lxxxvii]

Like the responsibility of the community as a whole, the responsibility of an individual *just* in virtue of her attitudes does not satisfy all four requirements for moral responsibility. But reading statements of blame as directed at the non-causal attitude holder is important for understanding the blame statements as a whole.

The Individual as Causal Agent

One reading of the responsibility ascriptions in question, then, is a pinning of blame on the shared attitude community itself. Another reading considers the individual as a non-causal attitude holder. Neither of these readings includes direct causal links between the subject and the harm. But "racists are responsible for racially motivated hate crimes" is a causal statement. So where is the causality?

The holding of a morally reckless attitude can be construed causally in terms of failure to act—the attitude-holder is failing to act to prevent herself from causing harm, should the proper situation arise. Say, for example, I am aware of my own homophobic attitudes. Many friends have told me that my beliefs are wrongheaded, and that I ought to rethink my stance, but I usually counter by saying, "it's not like I'm hurting anyone; I just think it's weird and unnatural and I don't want to see it." One day, I read a newspaper article about a hate crime whose target was two members of my local gay community. It was a crime of opportunity, not thought out or premeditated, but perpetrated by teenagers who happened to see two men holding hands in a park. I think, "that could have been me, I could have done that." At this point, the recklessness of my attitudinal orientation has hit me; it is only the case that my homophobia is not hurting anyone because I have not been in a situation in which my causing direct physical or material harm is likely. If, at this point, I make no effort to think further about the harmful nature of my attitude, I

am responsible for failing to engage in the kind of introspection that could deliver me from my moral recklessness.

However, to hold an individual responsible for an attitude is not the same as saying that an individual is responsible for membership in a shared attitude community. There is more to say about the individual members of a shared attitude community than that they hold an attitude that makes them blameworthy. It is their existence as a part of a community that is unusual; there are plenty of attitudes in the world that are not necessarily detrimental, although they could be, because there is no community sharing them. For example, I had a friend in high school who fervently believed that if we were to dump sterilizing agents into the world's Mountain Dew supply, we would be living in utopia within three generations. This was an attitude she had; if she had access to sterilizing agents and the world's supply of Mountain Dew, she would have acted upon her belief. (That is probably not really the case, but for the sake of this argument I will assume that she was genuinely a potential soft-drink terrorist.) However, she is, to my knowledge, the only person to hold her Mountain Dew attitude. Or if there are others with a similar attitude, they do not necessarily have access to each other. Because of her solitariness, she really posed little threat to the Mountain Dew drinkers of the world. Not only did she not have the means to spike the supply, she was not influencing anyone else, or creating a hostile environment for Mountain Dew drinkers, or contributing to an atmosphere of impunity for others wishing to take radical action against the Mountain Dew-drinking world. There is something to being a member of a shared attitude community beyond just possessing the shared attitude in question.

Although it cannot be the case that the individuals implicated in blame statements like the ones in the introduction are being held responsible *just* in virtue of their attitudes, in also cannot be the case that "the German people were responsible for the Holocaust" means that every German citizen, as an individual, is responsible for the actual murder of each person killed during the Third Reich. Most of the people belonging to the community in question never pulled a trigger or flipped a switch. To suggest that such an individual is responsible for those actions would not satisfy any of the four criteria for moral responsibility put forth in the first chapter, and

would be committing the kind of unethical barbarism against which H.D. Lewis warns.[lxxxviii] One ought not be held responsible for the actions of another.

Furthermore, it is not useful, for the purposes of explicating the blame statements at issue in this chapter, to go into too much detail about individuals who do, in fact, pull triggers or flip switches. The individual perpetrators of direct physical or material harm are certainly subject to blame, but their responsibility is assessed in individual, not collective, terms. In the case of direct material harm, then, only the person with the direct causal link can be considered morally responsible and blameworthy, and their actions are of interest in this chapter only insofar as they are products of the shared attitude community. Because responsibility for material harm is tied to such strict causality, the blaming of shared attitude communities is incoherent unless some other harm is being caused.

So if the harm is not the direct material or physical harm associated with harmful attitudes, what harm *is* at issue? It seems more likely that the harm in question is the perpuation of a community that may itself lead to harm. The community may create an atmosphere in which material harm is likely, or the existence of the community may lead to reasonable fear.

May subscribes to this description of harm in addition to his moral recklessness account. Shared attitude communities create an atmosphere in which material harm is likely, in that it is tacitly condoned, codified, embraced, or willfully ignored. As May writes: "The individual racist attitudes considered as an aggregate constitute a climate of attitude and disposition that increases the likelihood of a racially motivated harm."[lxxxix] A shared attitude community, and the individual members who contribute to the community's existence, are responsible for the harm created by the community itself. This harm is not the direct physical or material damage, but the atmosphere created by the existence of a number of people with the same harmful attitude.

Also, the existence of a community centered on a harmful attitude leads to reasonable fear, which is, as Linda Radzik argues, "a morally significant harm."[xc] Radzik explains that reasonable fear has two compontents. The first is epistemic: reasonable fear must be based on something more than pure fantasy. Conspiracy theories and paranoid delusions that a community

of people is out to harm you do not count as reasonable fears. The other component is moral: Reasonable fear must not be based on immoral prejudices or one's own wrongdoing. For example, during the Jim Crow era, a popular mythology led people to believe that all black people were cannibals. Thus, white people may have feared black people simply because of an immoral prejudice against black people. Radzik does not consider such fear reasonable. She also gives an example of a weaker group being oppressed by a stronger group. The stronger group may fear that the weaker group may rise up in rebellion, and that fear may be epistemically reasonable. However, there is no duty for the weaker group to respond to the fears of the stronger group, because the stronger group's fear is based on their own wrongdoing.[xci] The harm in question, then, is the perpetuation of a morally harmful community, one that makes material harm or reasonable fear commonplace.

There are many ways that an individual attitude-holder can be considered directly responsible for harm. Besides the obvious individual participation in hate crimes or other attitude-incited atrocities, an individual is causally tied to the existence and the perpetuation of a community that fosters harm and fear. The degree to which an individual participates in the shared attitude community may vary—a person who writes a letter-to-the-editor of her local newspaper arguing against holding an integrated prom is not responsible to the same degree as a Ku Klux Klan member. But, being causally connected to the harms of shared attitude community requires considerably less participation than actual perpetration of crime.

Putting on the 3-D Glasses

Returning to the blame statements proposed in the introduction: "the German people are to blame for the Holocaust," "as an American, I feel guilty for the war in Iraq," "the Rwandan Hutus are responsible for the genocide of Tutsis and moderate Hutus in 1994," and "racists are responsible for the legacy of oppression and violence against blacks in America," it is now clear that these grammatically simple sentences are much more complicated than they appear. The subject being held responsible is a shared attitude community—a group of people who are united

because they have an attitude in common. Unpacking the concept of a shared attitude community reveals that there are many kinds of agents and harms involved: the community itself, the individual attitude holders in absence of any causal connection to harm, the individual to the extent that she participates in the creation and perpetuation of the community, direct physical or material harm, the harm engendered by the existence of a community in which direct physical harm is likely, reasonable fear…each of these readings is a part of the broader picture of responsibility attributions to shared attitude communities.

The shared attitude community is not a singular bearer of moral responsibility, like the institution. Rather, it is a convenient way of speaking about a large and conglomerate group of individuals. Although it is not quite right to say that shared attitude communities are collectively morally responsible for certain harms, using the language of blame and responsibility to describe the actions of shared attitude communities is not without instrumental importance. In particular, blame statements made about attitude communities incite community members to ask important questions about their own roles in harm.

It is important to present the shared attitude community analytically, as well as synthetically; by looking at individual components of the community it is possible to see more clearly what actions can be asked of community members. For example, because the community itself is not a decision-making causer, it is not possible to expect it to change. Unlike an institution, which may act institutionally to bring about change in response to blame or praise, a shared attitude community cannot decide to act in a morally praiseworthy manner. There is no directive uniting a shared attitude community; it is incoherent to say "the racists decided today that they would cease their immoral actions." Furthermore, any such declaration would be tantamount to a disbanding of the shared attitude community: "the racists decided they would no longer be racist." These statements simply don't work. There is little sense in expecting the community itself to transform; the response to blame must be an individual one, taken in the context of a community that transcends individual aggregation.

The individual member of a community has, to borrow a term from Radzik, a "duty to respond."[xcii] Regardless of a person's degree of participation in the shared attitude community, every person has some obligation to react to harm. Or, to go back to the distinction between taking responsibility and holding responsible, one may take responsibility without being held responsible. Individuals may have duties to respond to the harms being allowed because of the existence of shared attitude communities; we may encourage them to take responsibility and change their own actions to try to discourage harmful attitudes or pay off their own moral debts.

[lxxiv] See, for example, Farid Abdel-Nour, "National Responsibility," *Political Theory* 31, no. 5 (October 2003), 719, http://0-links.jstor.org:693/sici?sici=0090-5917%28200310%2931%3A5%3C693%3ANR%3E2.0.CO%3B2-4.

[lxxv] Virginia Held, "Group Responsibility for Ethnic Conflict," *Journal of Ethics: An International Philosophical Review* 6, no. 2 (2002), 157–178.

[lxxvi] Larry May, *Sharing Responsibility* (Chicago: University of Chicago Press, 1992), 204.

[lxxvii] ibid., 46

[lxxviii] Aggregate and conglomerate distinction from Peter A. French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984).

[lxxix] May, *Sharing Responsibility* 48.

[lxxx] May, *Sharing Responsibility*, 49.

[lxxxi] ibid., 50.

[lxxxii] Linda Radzik, "Collective Responsibility and Duties to Respond," *Social Theory & Practice* 27 (July 2001), 459.

[lxxxiii] ibid., 458.

[lxxxiv] ibid., 459.

[lxxxv] ibid., 459.

[lxxxvi] ibid., 459.

[lxxxvii] For more on the distinction between holding responsible and taking responsibility, see Claudia Card, *The Unnatural Lottery* (Philadelphia: Temple University Press, 1996), chapter 2; and chapter 1 of May, *Sharing Responsibility*, 204.

[lxxxviii] H. D. Lewis, "Collective Responsibility," *Philosophy: The Journal of the Royal Institute of Philosophy* 23 (January 1948), 3–18.

[lxxxix] May, *Sharing Responsibility*, 47.

[xc] Radzik, *Collective Responsibility and Duties to Respond*, 467.

[xci] ibid., 467.

[xcii] ibid.

Conclusion

I began this paper with the question: What is collective moral responsibility, and why should you care? I end with a more specific question: what does it mean for collective responsibility to think in terms of institutions, situational collections, and shared attitude communities; and why do *I* care?

As I mentioned in the introduction, and reiterated throughout the paper, collective responsibility is a beleaguered notion. People have been deeply uneasy with the prospect of allowing non-human entities into the moral community, and doomsayers have associated the rise of collective responsibility ascriptions with the very end of ethics itself. Collective responsibility has typically been quite a lot to swallow. I have sought to break "collective" down into more manageable bites: institutions, situational collections, and shared attitude communities. By increasing the specificity of collective responsibility discussions, I hoped to add to the coherence of the theory as a whole. There are some kinds of groups that may legitimately and ethically be considered collective moral agents: institutions, for example, are capable of bearing moral responsibility. For other kinds of groups, collective responsibility is not the best way to understand the assignation of moral responsibility—situational collections are better described using models of shared cooperative action and shared responsibility. Finally, as is the case for shared attitude communities, some groups can only be understood as having characteristics that make them both like and unlike a collective. To arrive at that kind of hybrid conclusion, though, it was crucial to first have clear understandings of the collective responsibility of institutions and the non-collective responsibility of situational collections. Only having looked at collectively responsible and cooperatively responsible groups could such a description of shared attitude communities have worked.

Looking at shared attitude communities, then, is particularly useful for understanding the import of collective responsibility, broadly construed. Talk of institutions as collective moral agents and situational collections as non-collectively responsible is blunt and simplified. In theory, clear delineations between collective and non-collective moral responsibility are possible;

in practice, it would be naïve to expect such clarity. So the ultimate conclusion of the discussion of shared attitude communities, which blur the lines between collective and non-collective agency, is of particular value. Individuals ought to consider their actions within the context of collectives and communities. We ought to understand both that we are individual moral actors, and that we exist side-by-side with things like institutions and communities in the moral realm. We should be willing to consider institutions morally responsible, in order to hold them accountable for their actions. We should be just as willing to take responsibility for the harmful attitudes we hold or the benefits we receive in light of the harmful attitudes of others. In order to make sense of these holistic moral claims, I first had to deconstruct and analyze.

BIBLIOGRAPHY

Abdel-Nour, Farid. "National Responsibility." *Political Theory* 31, no. 5 (October 2003): 719, http://0-links.jstor.org:693/sici?sici=0090-5917%28200310%2931%3A5%3C693%3ANR%3E2.0.CO%3B2-4.

Aristotle. *Nicomachean Ethics*. Edited by Terence Irwin. Indianapolis, Ind.: Hackett Pub. Co., 1985.

Austin, J. L. "Three Ways of Spilling Ink." *Philosophical Review* 75, (October 1966): 427–440.

Bratman, Michael. "Shared Cooperative Activity." *Philosophical Review* 101, no. 2 (April 1992): 327–341.

Bratman, Michael. "Shared Intention." *Ethics* 104, no. 1, (October 1993): 97–113.

Broome, John. "Is Incommensurability Vagueness?" in *Incommensurability, Incomparability and Practical Reason*, edited by Ruth Chang. Cambridge, Mass.: Harvard University Press, 1997.

Card, Claudia. *The Unnatural Lottery*. Philadelphia: Temple University Press, 1996.

Chang, Ruth. "The Possibility of Parity." *Ethics: An International Journal of Social* 112, no. 4, (July 2002): 659–688.

Elworthy, Scilla and John Beyer, eds. *How Nuclear Weapons Decisions Are Made*. New York: St. Martin's Press, 1986.

Enron: The Smartest Guys in the Room. DVD. Directed by Alex Gibney. 2005.

Feinberg, Joel. "Collective Responsibility." *Journal of Philosophy* 65, (November, 1968): 674-687.

French, Peter A. "Principles of Responsibility, Shame, and the Corporation." in *Shame, Responsibility, and the Corporation*, edited by Hugh Curtler. New York: Haven Publications, 1986.

———. *Collective and Corporate Responsibility*. New York: Columbia University Press, 1984.

Held, Virginia. "Can a Random Collection of Individuals be Morally Responsible?" *Journal of Philosophy* 67, (July 1970): 471–480.

Held, Virginia. "Group Responsibility for Ethnic Conflict." *Journal of Ethics: An International Philosophical Review* 6, no. 2 (2002): 157–178.

Held, Virginia. "Moral Responsibility and Collective Action." In *Individual and Collective Responsibility: The Massacre at My Lai*, edited by Peter A. French. Cambridge, Mass.: Schenkman Pub. Co., 1972.

Kutz, Christopher. "Acting Together." *Philosophy and Phenomenological Research* 61, no. 1 (July 2000): 1–31, http://links.jstor.org/sici?sici=0031-8205%28200007%2961%3A1%3C1%3AAT%3E2.0.CO%3B2-%23.

Lewis, H. D. “Collective Responsibility.” *Philosophy: The Journal of the Royal Institute of Philosophy* 23, (January 1948): 3–18.

Mackie, J. L. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin Books, 1977.

May, Larry. *Sharing Responsibility*. Chicago: University of Chicago Press, 1992.

Narveson, Jan. “Collective Responsibility.” *Journal of Ethics: An International Philosophical Review* 6, no. 2 (2002): 179–198.

Radzik, Linda. “Collective Responsibility and Duties to Respond.” *Social Theory & Practice* 27, (July 2001): 455–471.

Sverdlik, Steven. “Collective Responsibility.” *Philosophical Studies* 51, (1987): 61.

Tuomela, Raimo. “We Will Do It: An Analysis of Group Intentions.” *Philosophy and Phenomenological Research* (June 1991): 249–277.

Velleman, J. David. “How to Share an Intention.” *Philosophy and Phenomenological Research* 57, no. 1 (March 1997): 29–50.