

INDIVIDUAL RESPONSIBILITY FOR COLLECTIVELY COMMITTED HARMS

by

SABA BAZARGAN

A Dissertation submitted to the
Graduate School-New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Professor Jeff McMahan

and approved by

New Brunswick, New Jersey

May 2010

ABSTRACT OF THE DISSERTATION

INDIVIDUAL RESPONSIBILITY FOR COLLECTIVELY COMMITTED HARMS

By SABA BAZARGAN

Dissertation Director:

Jeff McMahan

I develop an account that addresses the moral permissibility of and responsibility for participating in cooperative ventures in which each member contributes marginally to a substantial, aggregate harm. I argue that a basis for assessing the morally permissibility of and responsibility for participation in such collectives should not be limited to the intentions and the causal contributions of the participant, but should also include the *purpose* – what I call the ‘functional role’ of the individual in the collective venture. This account has the benefit of grounding a reductionist metaphysical analysis of cooperative collective action. After developing such an account, I apply it to a particular type of cooperative, collective activity, viz., warfare. The account of contributory moral responsibility I develop provides a basis by which to ascertain moral responsibility for marginal contributions by combatants participating in an unjust war.

Table of Contents

Title page.....	i
Abstract	ii
Table of Contents	iii
List of Tables.....	iv
1. Introduction	1
2. On ‘Responsibility’	9
3. Metaphysical Conditions.....	16
4. Metaphysical Account of Collectives.....	20
5. Metaphysical Conditions Revisited.....	56
6. Normative Conditions	66
7. An Account of Responsibility for Collective Action	85
8. War - A Test Case	181
Bibliography	208
Curriculum Vita	211

Lists of tables

Mereology of Warfare	189
Moral Relations of Elements in War 'W'	199

1. Introduction

Many of the world's worst harms are the result of what individuals do together. That is, many of the world's worst harms are the result of collective action. Yet there have been surprisingly few attempts to develop an account of individual responsibility for collective action. My goal is to develop such an account.

An account of individual responsibility for collective action requires a metaphysical account of collective action. Attempts by others to develop such an account can be categorized broadly into two types. One type claims while the other denies that collective actions in general are reducible to (non-trivial) properties of individual persons. The accounts of collective action that affirm this reducibility can be called "reductive accounts" of collective action. Accounts that deny the reducibility of collective action can be called "non-reductive accounts" of collective action. Reductive accounts of collective action have to their advantage ontological parsimony. Non-reductive accounts of collective action have to their advantage a supposed basis for our beliefs and practices that ostensibly reference collectives per se.

My account of individual responsibility for collective action will be based on a reductive account of collective action. I think that non-reductive accounts have remained recalcitrant partly because our lay-concept 'collective action' refers to too many kinds of events for which to provide a univocal analysis. For example, given an appropriate context, each of the following can, at least ostensibly, be the subject of a collective act: a corporation, all blondes, a government, a sports team, an army, a political party, a

religion, a nation, an ethnicity, a gender, a mob, an ad hoc committee, all cancer survivors, all tax payers, all car drivers, all sushi-eaters, a group playing soccer, three people pushing a car out of the snow, two friends walking together, every rich person in the past 200 years, etc.

It is not my intention to provide an analysis of our lay-concept 'collective action'. Nor will I present truth-conditions for ascriptions of the term "collective action". Rather, I will present a concept that carves out a particular species of collective action -- what I call "cooperative" collective action. Likewise, I will develop an account of individual responsibility for cooperative collective action specifically, rather than for both cooperative and non-cooperative collective action.

A cooperative collective act is an act performed by a cooperative collective qua cooperative collective. A non-cooperative collective differs from a cooperative collective in that the latter consists of individuals at least some of whom intentionally act together. At one remove, they act together in that their actions are coordinated; the individuals anticipate one another's decisions, and act accordingly. At another remove, they act together in the sense that they cooperate; the individuals coordinate their action in order to bring about a goal consisting in or caused by a combination of the contributions made by the individuals composing the collective. While the particular goals of various cooperative collectives differ widely, these collectives all share the feature of consisting of at least some individuals who not only coordinate but cooperate in order to bring about an end, even if this end is ill-defined, disjunctive, or vague. Examples of cooperative

collectives include: a football team, a political party, a multi-national corporation, the House of Lords, a small business, a group freeing a car from the snow, a platoon, friends walking together, an orchestra, an extracurricular club, and so on.

A non-cooperative collective, on the other hand, consists of individuals who unintentionally act together. Even though they do not intend to act together, the combination of their individual actions is a collective event. Two collectives can differ with respect to whether they are cooperative or non-cooperative depending solely on whether the individuals composing the collective have the appropriate intentions. For example, suppose each resident in an apartment building decorates his or her own window in order to contribute to lighting up the exterior of the entire building for the holidays. This collective act is a cooperative act. Contrast this case with a nearby example. Suppose each resident decorates his or her window with lights; the entire building is lit up as a result. But unlike the previous case, the residents do not intentionally contribute to lighting up the entire building. It is, instead, a side effect of each individual's intention to decorate her own apartment. This is so in spite of the fact that each resident of the building is aware that by lighting up her window she contributes to the collective act of lighting up the entire building. This collective act is non-cooperative.

I argue that cooperative collective acts are type-identical to a combination of individuals, individual attitudes, individual actions, and (non-trivial) individual relations. But more specifically, I argue that the properties to which cooperative collective action is reducible

necessarily include the functions of the individuals in the cooperative collective. Of course, there are many cooperative collectives in which some or all of its members have functions. But my claim is more radical than this; I argue that functional roles are fundamentally constitutive, not only of cooperative collectives with members who unobjectionably have functions, but of all cooperative collectives, including those consisting of members none of whom have apparent functions. I argue that this kind of reductive account fares better than competing accounts when confronted with a gauntlet of philosophical challenges.

In the section following the next, I lay out three metaphysical conditions that an account of cooperative collective action must satisfy. First, any theory of cooperative collective action must account for the fact that a cooperative collective can act even if some or most of its members do not contribute to that act. Likewise, any account of cooperative collective action must be compatible with the claim that we cannot infer that a cooperative collective committed an act from the fact that most (or even all) members of the collective together committed that act. Second, an account of cooperative collective action must explain how some collectives can apparently survive the replacement of all its members. And third, the account must elucidate the basis for the apparent obligations resulting from participation in cooperative collective action.

In section four, I present my account of cooperative collective action. On this account, a cooperative collective is a conjunction of functional roles that share one and the same object. A functional role has three parts -- a subject, a set of goal-oriented rules, and an

object. The subject is the person that occupies the functional role. Acting in accordance with the goal-oriented rules is productive with respect to a particular collective event, which is the functional role's object. For an individual to be a member of a collective, it is necessary for her to be the subject of a functional role. And to be the subject of a functional role, it is necessary for the individual to have committed herself to acting according to goal-oriented rules.

These criteria might seem too strong. It seems there are many cooperative collectives consisting of members who are not subjects of functional roles, in that they have committed themselves neither to contributing to a cooperative collective act, nor to acting according to goal-oriented rules. In sections 4.3 and 4.4 I argue otherwise. In so doing, I argue that if an individual agrees to act as part of a cooperative collective, and she knows or is in a position to know that the set of rules is goal-oriented with respect to a cooperative collective event, and if others who act according to rules with the same end are disposed to regard her as having a purpose in their cooperative collective, then she has a functional role. But precisely who are these 'others' who must be disposed to regard her as having a purpose in their cooperative collective? How many such individuals must there be? What if there is disagreement among them regarding whether she has a purpose? In section 4.5 I answer these questions by explaining how the relations among individuals in a cooperative collective ground functional roles.

In section five, I argue that the account I developed in the previous section satisfies the metaphysical conditions I presented in section three. In section six, I present the

normative conditions that an account of cooperative collective action must satisfy. I present two conditions, broadly construed, that any account of responsibility for cooperative collective action must satisfy. First, any account of responsibility for cooperative collective action must be compatible with the “Control Principle” according to which an individual is responsible for an event only if that individual had relevant control over the occurrence of that event. Second, if a cooperative collective commits a wrongful act, then there must be agents responsible for that act, even if only partially responsible. This condition can be called “The Responsibility Principle”.

After elucidating the Control and Responsibility principles and providing reasons for imposing them as conditions, I argue that these two conditions often pull in different directions when they are applied to cooperative collective acts; satisfying one principle seems to preclude satisfying the other. To resolve this dilemma, I distinguish between ‘primary’ responsibility and ‘secondary’ responsibility, where the former is responsibility for what one does, while the latter is responsibility for what others do. I argue that the Control Principle is incompatible with secondary responsibility; as a result, I argue that the Control Principle should be weakened so that only the basis of responsibility must be an event under the control of the responsible agent.

In section seven I develop an account of individual responsibility for cooperative collective action that builds upon the account of cooperative collective action developed in section four. The basis of an individual’s responsibility for a collective act is the individual member’s decision to become part of the cooperative collective -- i.e., to take

on a functional role. By agreeing to do so, the individual bears secondary responsibility for what the cooperative collective does, over and above the primary responsibility incurred by the difference her contribution makes. I argue that this account of individual responsibility is compatible with the weakened version of the Control Principle.

But why believe that voluntarily joining a cooperative collective -- i.e., voluntarily taking on a functional role -- is a basis for responsibility? I will argue that joining a cooperative collective by filling a functional role constitutes an implicit agreement which serves as a basis for secondary responsibility, even if the object of responsibility is not causally or counterfactually dependent on the member's implicit agreement.

It is not the case, however, that each member of a cooperative collective is fully responsible for every collectively committed act in a cooperative collective. In sections 7.4 through 7.6, I describe various ways that responsibility for collectively committed acts in a cooperative collective can be limited. At the end of section six I present a highly generalized but detailed picture of how responsibility is disseminated in nested cooperative collectives -- i.e., cooperative collectives that themselves include cooperative collectives.

In the final section, I apply the account of individual responsibility for cooperative collective action to a paradigm of cooperative collective action: war. Specifically, I will be concerned with the responsibility of individual combatants for acts committed by the military of a country waging a war. To this end, I implement the account of responsibility

in nested cooperative collectives, developed in the previous section. An intermediate conclusion will be that a just war cannot involve the pursuit of unjust subsidiary aims. Ultimately, I argue that it is often impermissible for a combatant to participate in an unjust war, even if her participation is limited solely to a just aim in that war.

A note on terminology: I will use the term ‘collective’ from here on to refer to cooperative collectives. Also, I will use lowercase letters to refer to individuals, uppercase letters to refer to collectives, and Greek letters to refer to acts.

2. On ‘Responsibility’

Why develop an account of individual responsibility for collectively committed harms?

To answer this, it is first necessary to clarify what I mean by the term “responsibility”.

There are multiple uses of “responsibility”; I will briefly describe each, after which I will clarify how I will be using “responsibility” here.

Responsibility-as-Causality. On this sense of ‘responsibility’, as the name suggests, an individual is responsible for ϕ just in case that individual causes ϕ . As a result, responsibility-as-causality is attributable to non-agents. For example, we might claim that lightening was responsible for a forest fire. In this sense of ‘responsibility’, the term “causally responsible” is pleonastic. This sense of ‘responsibility’ is both too weak and too strong for my purposes here. It is too weak in that the sense of ‘responsibility’ I have in mind is applicable only to agents. Responsibility-as-causality is also too strong, in that I want to leave open the possibility of coherently and correctly attributing responsibility to agents for a harm that they did not cause.

Responsibility-as-Attributability. This sense of ‘responsibility’ is intuitively compelling yet difficult to specify. Put metaphorically, ϕ is attributable to x in that x is the author or owner of ϕ . Clearly, though, more needs to be said, as actions are neither the sorts of things that can be owned, nor are they the sorts of things that can be authored. In paradigmatic instances of responsibility-as-attributability, x causes ϕ -- but causation is

not sufficient for responsibility-as-attributability. Instead, this sort of responsibility is perhaps best elucidated by considering its moral role. While responsibility-as-attributability is not itself a moral property, it satisfies a condition or acts as a basis for the attribution of certain moral properties, such as those necessary for bearing moral liability or culpability. (Responsibility-as-attributability might also be a basis for the attribution of non-moral but normatively relevant properties in certain cases such as the property of having provided a service). In any case, responsibility-as-attributability, though it is closely related to the sense of ‘responsibility’ that I have in mind, is too inchoate to itself serve as this sense of ‘responsibility’.

Responsibility-as-Obligation. Suppose ϕ is a bad event that x has not committed. Moreover, x is not liable to anyone for the occurrence of x . In effect, x is a bystander with respect to ϕ . However, x is in a position to respond to ϕ in a way that imposes a small cost on x , but which significantly lessens a harm or a wrong caused by ϕ . Suppose that it would be morally wrong for x to refuse to respond to ϕ . However, x cannot be permissibly forced to so respond, because (or more aptly, *in that*) x is not morally liable to anyone for the occurrence of ϕ .

For example, suppose there is a devastating car accident; y , a victim in the accident, is trapped in a burning car, and will die unless the only bystander, x , assists y . Suppose that for x to do so would impose little to no cost on x . In such a case, one might argue that x has a moral obligation to respond by assisting y , provided that doing so will prevent a

great harm at little cost to x. We might construe this obligation to respond as one of the senses of ‘responsibility’. Here, the object of responsibility is not a harm that has occurred, but a harm that x ought to prevent from occurring. Thus, this sense of responsibility is significantly distinct from the others -- x is responsible for helping the victim, even though she did not cause the accident, is not morally liable for the accident, and is not to blame for the accident. To fully distinguish this sense of ‘responsibility’ from moral liability, we might claim that x, in the present example, cannot be permissibly forced to help y, even though x is morally required to help y. Should x fail to help, she is at worst the appropriate subject of contempt, dismay, moral disgust, etc.

This kind of moral responsibility -- responsibility-as-obligation -- is important and interesting, but is not the type with which I am concerned here. I will focus on stronger versions of moral responsibility, namely those which assign moral liability or culpability to the responsible party.

Responsibility-as-Blame. Perhaps the most common way to think of responsibility when the subject of responsibility is an agent, is by construing it as blame. On this sense, an individual is responsible for ϕ just in case she is to blame for ϕ . The sense of ‘responsibility’ that I will use will include responsibility-as-blame, but will not be limited to it.

Responsibility-as-Liability. When responsibility is construed as moral liability, x can be subjected to or forced to perform an act, ψ , that remedies or prevents a harm, ϕ . I leave open the possibility that moral liability can be strict; thus responsibility-as-liability for ϕ does not entail blame for ϕ . That is to say, responsibility-as-liability does not entail responsibility-as-blame. I will argue that x need not have caused ϕ in order for x to be liable for ϕ . Thus responsibility-as-liability will not entail responsibility-as-causality. .

I will focus on two senses of ‘responsibility’: responsibility-as-liability and responsibility-as-blame. Unless I indicate otherwise, I will use “responsibility” disjunctively to refer to either responsibility-as-liability or responsibility-as-blame.¹ I will argue that a member of a collective might bear both responsibility-as-liability and responsibility-as-blame for an act the collective commits; alternatively, she might bear only responsibility-as-liability for the act the collective commits.

¹ Using “responsibility” in this way might seem to constitute a sort of semantic gerrymandering, in that I have by stipulation disjoined two concepts, i.e., liability and culpability. But this stipulation is not inappropriate. Both responsibility-as-liability and responsibility-as-blame fall under the concept ‘responsibility’. Put differently, the above senses of ‘responsibility’ are not disambiguations of the term “responsibility”. Rather, they are different senses of the concept ‘responsibility’. This is not a trivial distinction. For instance, we might disambiguate a use of the term “bank” by clarifying whether it refers to a financial institution or the side of a river. These two possible uses of “bank” do not constitute different senses of a concept ‘bank’. Rather, the two possible uses of “bank” involve two concepts which are not otherwise related (etymology aside). The term “responsibility”, on the other hand, is not ambiguous. Presumably, there are various fine-grained concepts that have family-resemblances to one another, all of which fall under the more coarse-grained concept ‘responsibility’. The disjunctive use of responsibility involves bracketing off various senses of ‘responsibility’. What remains is itself a sense of ‘responsibility’ since what remains falls under the concept ‘responsibility’. Of course, in doing Normative Ethics, we might discover that some senses of ‘responsibility’ refer to properties that are not ethically relevant. This would provide a reason to use ‘responsibility’ slightly differently from how it is actually used. But my disjunctive use of “responsibility” should not itself be taken as an implicit claim regarding how we ought to use the term “responsibility”, or how we ought to think of the concept ‘responsibility’.

In the discussion of responsibility for collectively committed harms, I will often claim that an individual bears *partial* responsibility for a harm. More generally, I will often claim that responsibility comes in degrees. Given the disjunctive use of “responsibility”, what it means to say that responsibility is “partial” depends on whether “responsibility” refers to moral liability or blame.

The claim that x bears partial *blame* for a harm ϕ , or that x bears a certain degree of blame for that harm, means that the blame itself, rather than the object of the blame, is limited. This has implications for, among other things, punishment of the blameworthy individual. In paradigm cases, if an individual is blameworthy for a harm, then this is grounds (though perhaps not sufficient grounds) for punishing the individual. The degree of punishment depends on the degree of blame for which the individual is punishable. If an individual bears more rather than less blame for a harm, and if that blame is a basis for punishment, then, *ceteris paribus*, the individual may be punished more rather than less severely.

Contrast degrees of blame with degrees of moral liability. If we take ‘moral liability’ as a three-place relation in which a subject is liable to a cost for an event ϕ , then it is the *cost* that is less than full if the liability is less than full, rather than the moral liability itself.

The amount or degree of compensation that the individual is morally liable to provide depends, of course, on the harm for which the morally liable individual is morally liable; but it also depends on the degree to which the morally liable individual bears moral liability for that harm. Suppose x bears less than full, or partial moral liability for a harm

ϕ . Since x is morally liable, she is morally required to bear the cost of a harm or wrong ϕ that a victim suffers (or is required to bear the cost of preventing that harm). But since x is not fully liable, she is not morally required to bear the entire cost of ϕ . So, for example, if an individual is partially liable for a harm that causes 1000 dollars in damages, and the individual is morally liable to provide compensation, then the amount that the morally liable individual must provide -- some percentage of 1000 dollars -- depends, in part, on the degree of moral liability that individual bears for the harm. If the morally liable individual bears a small degree of moral liability for the harm, then the individual will, generally, have to provide very little compensation for that harm.

It is more difficult to quantify partial moral liability when the liability in question is to preventive injury. In such cases, it is permissible to harm an individual if doing so will in turn prevent a harm for which that individual would be morally liable. When a large collective commits a harmful act, it is unlikely that every member who bears some moral liability for the ensuing harm will be in a position to prevent that harm. But if a member were in such a position, in that harming the member would prevent the collective harm, then the amount or degree of preventive harm to which the member is morally liable depends on the degree of moral liability that the member bears for that harm. If the degree of moral liability is low, then the amount or degree of preventive harm to which the member is liable is also low. So, for example, if kicking the morally liable member were the only way to prevent a substantial collectively committed harm, then doing so would be permissible in that the small degree of moral liability the member bears is enough to make him liable to preventive kicking. (Of course, if the collectively

committed harm is large enough, then kicking a *non*-member would be permissible if doing so would prevent the harm; but unlike the member who bears some responsibility for the collective harm, the non-member could permissibly demand compensation for the kicking since she bears no responsibility for the collectively committed harm). If the member bears only slight moral liability for the collectively committed harm, then severely injuring or killing the member in order to prevent that harm is likely to be impermissible, even if this is the only way to prevent the collectively committed harm. (Again, this assumes that the collectively committed harm is not so egregious that it would permit injuring or killing a non-member in order to prevent the harm).

The reasons why responsibility for a harm is, in general, important have partly been answered by clarifying what I mean by 'responsibility'. Responsibility is important because blame and moral liability are important. More specifically, responsibility is important because if an individual is responsible for a harm, then that individual may very well be, among other things, punishable for a harm or liable to compensation, repair, or preventive injury with respect to that harm. An account of responsibility for collectively committed harms allows us to determine when individuals are punishable, morally liable to compensation, preventive injury, etc. for collectively committed harms. Considering that many of the world's worst harms are caused by collectives, an account that provides the conditions of responsibility for such harms is eminently useful.

3. Metaphysical Conditions for an Account of Collectives

Before presenting an account of responsibility for collective action, it is necessary -- or at least helpful -- to present an account of the metaphysics of collective action. Any such account must satisfy many conditions. Here I present those that have proven especially challenging for reductive theories of collective action, of which mine is an example.

After doing so, I will develop an account of collective action that satisfies the conditions that I present in this section.

3.1. NON-COMPOSITIONALITY OF COLLECTIVES

The relationship between those who constitute a collective and the actions of the collective itself has two characteristics with which any reductive theory of collective action must be compatible.

Firstly, any theory of collective action must account for the fact that collectives can act even if some or most of their members do not contribute to that act. In fact, some collectives can act even if only one member acts. For example, the CEO of a corporation, or the President of the United States, has discretionary powers used to implement ad hoc policy without vetting by the board of directors or by Congress. In virtue of the fact that the President or CEO committed the act, the collective itself has also committed that act. We can call this characteristic, in which a collective acts even though some or most of its members do not, the “supra-compositionality” of collective action.

Secondly, any account of collective action must be compatible with the claim that we cannot infer that a collective committed an act from the fact that most (or even all) members of the collective together committed that act. As an example, suppose the Board of Directors of corporation X votes to reduce the quarterly dividends. After they have adjourned their meeting, and after the work-day has ended, they decide to eat together, and resultantly vote to dine at a sashimi restaurant. Only the first act of voting resulted in a collective event that had X as its subject. But this is not so for the second act of voting, even though it was performed by the same people. So though it is true that if X commits a collective act, then some members of X contributed to a collective act, it is not the case that if members of X contribute to a collective act, then X commits a collective act. Call this characteristic the ‘sub-compositionality’ of collective action.

Any theory of collective action must be compatible with the super- and sub-compositionality of collective action. The alternative is a commitment to summative or quasi-summative accounts of collective action, which has been adequately criticized by Margaret Gilbert (1989).

3.2. NON-EXTENTIONAL IDENTITY CONDITIONS

Our attitudes and interests largely determine our taxonomy of objects. For example, our interests and attitudes determine the reference-fixing conditions associated with the concept ‘chair’. We could have eschewed our actual concept in favor of one that encompasses both chairs and stools. If we had, the physical structure of the world would

not be different (Idealism aside). Although the identity conditions of the concept 'chair' depend on our attitudes, our attitudes are not among its identity conditions.

However, there are certain kinds of objects the existence of which logically depends on our attitudes. Cooperative collectives are among them, in that certain attitudes are among the identity conditions of such collectives. Without the relevant attitudes, the collective would literally cease to exist; there would only be various individuals. In contrast, if the attitudes that fix the reference of 'chair' ceased to exist, this would not cause chairs to disappear. The challenge, then, is to explain what attitudes are among the identity conditions of collectives, and *whose* attitudes are those upon which the collective logically depend.

It is not merely the identity of a collective at a time that needs explaining. In certain types of collectives (corporations, football teams, nations), the collective's extentional parts underdetermine its identity in the additional sense that the collective can survive a replacement of all its members. Any account of collective action must provide perduring/enduring conditions that allow particular kinds of collectives to survive the replacement or loss of their members.

3.3. NORMATIVITY OF COLLECTIVE ACTION

Reductive accounts of collective action, such as John Searle's (1990) and Michael Bratman's (1993) have been criticized on the grounds that they fail to explain the

normativity unique to individuals acting qua members of a collective. Margaret Gilbert (1989) has argued that membership in a collective involves an intention to act with others. Such an intention entails implicitly forming certain commitments to other members of the collective, who are justified in expecting compliance with these commitments. A failure to act on this obligation invites appropriate second-personal attitudes of rebuke and resentment. Consider Gilbert's oft-noted example of walking together (*supra*). Suppose that you and I form and act on respective intentions to walk together. During the walk, I change direction and walk away from you. Given our intentions to walk together, my action does not simply violate norms of practical rationality -- it also violates a tacit or implicit commitment between you and me. The commitments implicit in collective action provide a reason to act; to the extent that these reasons are motivating, they can explain individual behavior in collectives. A challenge to theories of collective action is to explain the source of the normativity implicit in collective action. To this end, Gilbert resorts to a non-reductive account of collectivity, underpinned by a technical notion of 'joint commitments' (*supra* 198). But I will present a reductive account of collective action that accounts for the obligations that members of a collective have to one another.

4. Metaphysical Account of Collectives

In the last section, I presented conditions that an account of the metaphysics of collective action must satisfy. In this section I present my account of collective action. I will explicate and defend this account in section five. In laying out my account of cooperative collectives in the following sections, I will provide few, if any, arguments in its favor. I am building a system from the ground up, rather than patching together an account by appealing to thought experiments and intuitions piece-meal. This is not to say that I ignore the plausibility of the account's practical consequences. But since our intuitions regarding individual responsibility for collective action often rely on incoherent or naive theories of collective action, it is important to attenuate the degree to which we rely on such intuitions in developing a metaphysical account of collective action. What will ultimately speak in favor of the account I present will be its ability to solve well-recognized problems in both the metaphysics and the ethics of collective action. As a result, I will develop an account of individual responsibility for collective action not simply by extrapolating an account using common-sense intuitions as data points, but instead chiefly by considering its cohesion with the metaphysics of collective action. Of course, the ethics of collective action cannot simply be "read off" the metaphysics of collectives; the latter underdetermines the former. The metaphysical account I provide, however, guides the investigation by constraining possible accounts of the ethics of collective action. The metaphysics carves out a wide valley; precisely which path we take is determined by an appeal to ethical intuitions, the ethical foundations of relevant legal policy, as well as various normative and ethical theories the application of which are especially cogent. As is the case with philosophical accounts that result from system

building, the account will occasionally yield consequences that conflict with intuitions. In such cases, the account is answerable to these intuitions, in that, at the very least, a credible error theory must be provided when the account deviates from our intuitions.

Ultimately, that the account provides a univocal solution to challenges in both metaphysics and ethics is what will yield the strongest argument in its favor. However, because in the next few sections I do not present arguments in favor of the account, the account may seem implausible. More specifically, the various claims regarding how we ought to think of collectives might seem too strong -- they might seem to rule out a host of candidates that are paradigm examples of collectives. However, in subsequent sections, I will show that the plasticity of the account silences earlier worries.

Put very briefly, on my account, a collective consists of individuals who act according to a functional role with one and the same object. A collective act is an event resulting from or consisting in a combination of acts committed by individuals acting in accordance with their functional roles.

4.1. FUNCTIONAL ROLES

On my account, a collective is a conjunction of *functional roles*. In what follows, I elucidate the concept ‘functional role’ after which I explain how members of a collective come to have functional roles.

A functional role, as understood here, has three parts. First, the subject of a functional role is the person that occupies the functional role. But for a person to do so, she must satisfy certain eligibility conditions. The eligibility conditions for satisfying the subject of a functional role should not be confused with the actual person that fills the functional role. I will refer to the latter when I use the term ‘subject’ or ‘subject-particular’. Second, the imperatives that compose the content of the functional role itself consist of goal-oriented rules that are (designed to be) productive with respect to a particular collective event. This collective event is the functional role’s object, which is the third part of a functional role.

The object of a functional role, the eligibility conditions of the role’s subject, and the goal-oriented rules of the functional role, are all determined by implicit or explicit agreements and commitments between a candidate subject, and members of the collective. (Later I will say more about how such agreements and commitments are formed). The goal-oriented rules constitutive of the functional role derive their normative force partly from the implicit or explicit commitment to abide by them.

Counter-intuitively, it is easier to characterize highly sophisticated functional roles, such as those that take the form of professional employment in a corporation, government, or other formal collective, than it is to characterize simple, inchoate functional roles.

Sophisticated roles, such as formal employment, typically involves taking a ‘position’, which largely (though often misleadingly) provides a ready-made characterization of that individual’s functional role in the collective. It is harder to characterize functional roles

that are exceedingly simple, associated with collectives lasting hours or even seconds, in which there is no explicit job or assignment of tasks. In such collectives, the object of the functional role, the eligibility conditions of the subject, and the content of the functional role, are all agreed upon the same way that implicit promises are reached. Scanlon's Expectationalist theory of promises (1990) can be used to explain how implicit agreements occur. (I will return to this issue when discussing the normativity of collective action).

With the concept of a functional role, it is possible to state in general 1) the conditions for membership in a collective and 2) the identity conditions of a collective. (Recall that by 'collectives' I mean 'cooperative collectives' only). For s to be a member of a collective, it is necessary for s to be the subject of a functional role. To be the subject of a functional role, it is necessary for s to have committed herself to acting according to goal-oriented rules. A collective is, in turn, a conjunction of functional roles that share one and the same object.

In the next section, I argue that positing functional roles has explanatory power with respect to our decisions to join collectives. This explanatory power does not alone provide adequate justification for thinking of collectives in terms of functional roles -- but it is a step in that direction.

4.2. CONTRIBUTORY INTENTIONS AND FUNCTIONAL ROLES

To show how the concepts of ‘functional role’ can help explain why we join collectives, it will first be helpful to elucidate (as I use it here) the concept ‘contributory intention’.

A contributory intention is an intention to contribute to an event, Σ , which consists partly of individual contributions made by others. That is, a contributory intention is an individual intention to commit an act for promoting Σ . The reason for contributing to Σ can be instrumental or otherwise.

An intention to contribute to Σ cannot be redescribed, in Davidsonian fashion, as an intention to commit Σ . This redescription would be illicit, because Σ is defined as an event that consists partly in the contributions made by other individuals. For example, a typical foot-soldier might intend to contribute to winning a war. But she cannot coherently intend to win the war. A voter might intend to contribute to the election of a candidate. But the voter cannot coherently intend to elect that candidate. Claiming otherwise strains conceptual coherence. It is precisely the fact that the object of intention is achievable by the agent that distinguishes intending from wishing or merely wanting. (For more on this point, see David Velleman’s (1997) arguments against Michael Bratman’s account of collective action).

It might be argued, however, that if an individual’s contribution is necessary for the achievement of Σ , then her intention to contribute can indeed be redescribed as an intention to Σ . Suppose that a collective consists of members m_a through m_z . A contribution made by m_z is necessary for the achievement of Σ . We can even suppose

that the contributions of the collective's members are serially ordered, so that ma contributes first, and mz contributes last. Without mz, Σ will not occur. But even in this case, when it is mz's turn to contribute, she cannot coherently intend to commit Σ . This is because intentionally committing an act entails a significant degree of authorship over that act. The act is attributable to the agent (recall the 'responsibility-as-attributability' sense of 'responsibility'). A first-personal intention on the part mz to achieve Σ entails authorship over Σ , of a kind or degree incompatible with the fact that there were other contributors as well. The upshot is that even when an individual's contribution is necessary for a collective event, that individual can at best intend to contribute to that event.

Though an individual with a coherent contributory intention cannot intend to commit Σ , it is possible for her to determine whether she commits an act that contributes to Σ . She can satisfy this intention to promote Σ precisely by committing a contributory act. Here it is up to the individual -- not others with whom she acts -- whether she achieves her contributory intention. This is because the satisfaction of an intention to contribute to Σ does not itself require Σ to subsequently occur.

Clearly, having a contributory intention -- that is, intending to contribute to a collective act -- is not enough for the contributor to be a member of a collective. Suppose an individual, mi , applies for membership in a club, but is rejected. Nonetheless, mi admires the club and its goals; as a result, he intentionally donates money to the club in order to further their cause. Certainly this does not make him a member of the club. Indeed, as a

member, m_i might contribute *less* than he would as a non-member. It is in virtue of adopting a functional role, not in virtue of how much he contributes, that m_i is a member of a collective.

Shared contributory intentions are neither necessary nor sufficient for membership in a collective; but they often *explain* implicit or explicit agreements to act according to functional roles. There are chiefly two ways a functional role can be related to a contributory intention. A pre-existing contributory intention can result in adopting a functional role. Or, a functional role can yield a contributory intention. Hence, a functional role can explain, or be explained by contributory intentions. I will start with the latter.

If an individual already possesses a contributory intention, then the individual has a *prima facie* motivating reason to act accordingly. Why would an individual implicitly or explicitly agree to act in accordance with a functional role? Functional roles are designed to be productive with respect to their object, a collective act. If an individual's contributory intention has as its object the same event to which the functional role is related, then that individual has a *protanto* instrumental reason to act according to a functional role. Acting in accordance with the functional role has instrumental value for the individual.

Alternatively, suppose an individual has no particular intention to contribute to a collective act. She agrees to act according to a functional role, not to satisfy a

contributory intention, but in exchange for something else of value, such as currency. In such a case, there is again an instrumental reason to act according to a functional role -- but not because it satisfies a contributory intention, but rather because it satisfies a personal end.

While the productivity of a functional role with respect to its object might help explain why a member might be motivated to contribute to the collective act, there seem to be many collectives which consist of individuals who have no goal-oriented rules for contributing to an end. In such collectives, even if individuals have an incentive to contribute, there are no functional roles by which to abide. Other collectives might not only seem to lack functional roles, but might also seem to have no 'end' to which to contribute. That is, it seems that while *some* cooperative collectives might include functional roles, others do not. I claim otherwise in the next section.

4.3. INFORMAL AND INCHOATE COOPERATIVE COLLECTIVES

In developing an account of the metaphysics of collective action, I have relied heavily on 'functional roles'. But it seems there are many collectives consisting of members who are not subjects of functional roles, and who have committed themselves neither to contributing to the object of their role, nor to acting according to goal-oriented rules. For example, an individual who joins a political party can deny that she has agreed to act according to any goal-oriented rules. For membership in this particular party, there is no explicit 'role' for individual members. Similarly, an individual who joins a running club

might deny that she is the subject of a functional role. The club consists of individuals who gather every week to run together -- no one is assigned any particular roles. The criteria I have provided for membership in a collective seem too strong, in that they do not seem to apply to many collectives.

But this is not so. Voluntarily joining any collective entails taking on a role which requires acting according to rules the adherence to which contributes to a collective act. However, in many collectives, the elements of a functional role (the object and the goal-oriented rules) are inchoate and informally established.

Consider *r*, who joins the Runner Club, which consists of individuals who run together for recreational and health purposes every Sunday. *r* does indeed have a functional role in virtue of agreeing to join the collective. Her role is the same as everyone else's, which is to run on Sunday with other members. If she fails to adhere to the rules in which her role consists -- e.g., if she doesn't show up, or if she shows up but bikes instead of runs -- then she has violated the justified expectation of others, from which the rules constituting her role are derived. (I will say more about the relevant expectation shortly). If *r* antecedently warns the other members that she will be unable to join them only on certain weekends, then the role and the goal-oriented rules in which her membership consists is affected accordingly.

There are collectives with roles that are more unremarkable than that of the Runner's Club. Consider a political party, membership in which requires nothing from its members

-- no fees, no meetings, no voting etc. The only purpose of this political party is to retain ideological solidarity. Each member, then, has a functional role, consisting in contributing to ideological solidarity, simply by following the rule consisting in the requirement of possessing the appropriate ideological views. If a member joined the party while holding views contrary to those for which the party stands, or if an existing member's view oppose those of the party, then her continued membership in the collective violates a commitment. She violates (even if no one discovers the violation) the justified expectation of others, from which the rules constituting her role are derived.

An even more vacuous cooperative collective is possible. Suppose a political party does not require even an interest in a particular set of views. The senior party members attempt to gain political influence simply by padding the collective's membership, (much as the Nazi party did in its early days) by enrolling as many members as possible. In this case, the role an individual member plays consists in contributing to the influence of the party by padding its membership, which requires nothing of the member save perhaps a disposition to affirm that membership when asked. A collective that literally has no purpose, and has literally nothing required or suggested of its members, can scarcely at all be called a cooperative collective.

It might be complained that this method of explaining away apparent counterexamples to my view involves shoe-horning the more troublesome counter-examples into my account, by arguing that the relevant functional roles 'are there', if only very hard to see. I hope this is not what I am doing. My intention is to provide a useful and revealing account of

what intentionally acting together fundamentally involves. I claim that my account is revealing of the metaphysical and normative structure of even these less cohesive collectives. This will, hopefully, become more evident as I continue.

4.4. THE OPACITY OF INTENTIONAL CONTRIBUTIONS

I have claimed that a collective is a conjunction of functional roles that share one and the same object, and that an individual is a member of a collective just in case she is the subject of a functional role. And an individual is the subject of a functional role just in case she commits herself to acting according to goal-oriented rules productive with respect to the object of the functional role. This is true, as I have argued in the previous section, even of collectives that do not ostensibly include functional roles.

But the claim that all members of every collective commit themselves to acting according to a functional role, might seem implausible. It seems that making such a commitment is not necessary in order to join paradigm instances of cooperative collectives, such as corporations, armies, clubs, political parties, etc. That is, it seems there are many collectives consisting of members who have not committed themselves to acting according to a functional role. Here I argue otherwise.

One way for a candidate member to commit herself to acting according to a functional role, is to agree to adopt that functional role. *Agreeing* to adopt a functional role (or ‘agreeing to a functional role’ for short) does not require an agreement that has in the

description of its content a term referring to a functional role. Neither must the candidate member have the concept of a functional role, in order to adopt a functional role. Rather, agreeing to a functional role consists in agreeing to:

1. contribute to the appropriate collective event
2. act according to goal-oriented rules

An individual agrees to a functional role just in case that individual agrees to 1 and 2. Call the claim according to which membership in a collective entails a commitment to contribute to the appropriate collective event, ‘the contributory commitment principle’ (CCP). And call the claim according to which membership in a collective entails a commitment to acting according to goal-oriented rules, ‘the rule commitment principle’ (RCP).

Identifying the act of agreeing to a functional role with the act of agreeing to (1) and (2), only pushes the problem back. Many collectives have members who have not agreed to contribute to a particular collective event. And most collectives consist of members who have not agreed to act according to goal-oriented rules. For example, the employees of a business might agree to no more than to perform certain tasks regularly. That these tasks also contribute to a collective event, or that following these tasks entail conforming to goal-oriented rules, are rarely salient to members of such collectives.

In the next sub-section, I lay out prima reasons to believe that the contributory commitment principle and the rule commitment principle are mistaken, after which I defend both against these criticisms. I argue that members of a collective often agree to take on a functional role without intending to do so. I also argue that members who do not agree, intentionally or otherwise, to contribute to the object of a functional role or to acting according to goal-oriented rules, have still committed themselves to contributing to that object by acting according to goal-oriented rules, and as a result have taken on a functional role.

4.4.1. PRIMA FACIE ARGUMENTS AGAINST CCP AND RCP

Here I will lay out prima reasons to believe that the contributory commitment principle and the rule commitment principle are mistaken. Typically, a good explanation for why an individual would agree to contribute to the occurrence of some event by acting according to goal-oriented rules, is that the individual desires the occurrence of that event. But there are many collectives with members who have no particular desire for the occurrence of the object of their functional role.

For example, suppose that mi enlists in the army and fights in a war. Though mi follows orders and thus contributes to the war's aims, he has no wish for the war's success.

Rather, his operative intention, when joining the armed forces, was to commit an act that would make him seem courageous to his friends and family at home. Joining the armed forces, and subsequently following orders, is instrumentally valuable to mi with respect

to the desired end of being perceived as courageous. The success of the war's aim is not valuable to mi . Suppose mi knows that he could contribute to the war's aims more effectively by staying at home and working in a lab. This fact does not provide mi with any motivating reason to do so, since mi is not concerned with the success of the war's aims. But mi does have a motivating reason to act as a soldier by following orders, since so doing is instrumental to his end, viz., being perceived as courageous by others.

The upshot is that, even absent a desire for Σ , mi can still be a member of a collective which commits Σ . It would be absurd to claim that, because mi does not desire the success of the war, he is not a member of the armed forces. There are countless collectives with members who care little whether the collective act to which they contribute occurs. In such cases, a desire for the occurrence of the object of their contribution cannot be plausibly attributed to such members. For example, individuals might join a collective intending only to get paid for performing a task, or to appear important to their neighbors, or to improve their CVs. Such an individual contributes to the appropriate collective event, since the contributory acts are also instrumental to achieving the individual's private ends. To be a member, it is enough that they agree to perform tasks required of members, the commission of which contributes to a collective act.

What explains the apparent implausibility of the contributory commitment principle and the rule-commitment principle, is that agreement is not closed under causal implication. An individual might agree to commit certain tasks, without agreeing to contribute to Σ ,

even if the individual knows that committing those tasks contributes to Σ . More generally, the fact that ϕ will cause ψ does not itself allow us to infer that *agreeing* to ϕ entails *agreeing* to ψ , even if the agent knows that ϕ will cause ψ .

It might be argued, in response, that all members, whether they are non-committal about the occurrence of the collective act Σ or whether they desire the non-occurrence of Σ , *do* indeed intend to contribute to Σ ; it is simply that their motivating reason to so contribute is not provided by a desire for the success of Σ . Rather, it is provided by a desire for the private end to which contributing is instrumental. To use an earlier example, contributing to the success of the war's aims is instrumental to m_i 's private end of appearing courageous -- so it might seem appropriate to attribute to m_i the instrumental intention to contribute to the war's aims. If so, then m_i would have had a motivating reason to agree to contribute to the war's aims -- even if he has no particular desire for the success of those aims.

But in such a case, there is little reason to regard the description of m_i 's intention as including a reference to a contribution to Σ . A member of a collective might agree to act in certain ways that, as far as she is concerned, contributes to the object of her functional role merely as a side-effect, rather than intentionally. Agreeing to commit a certain set of acts ϕ is not the same as agreeing to contribute to Σ , even if ϕ is instrumental to Σ .

To see this more clearly, consider an example of an individual who hopes that the collective of whom she is a part *fails* to meet its goals. This is in contrast to the above

example in which m_i is non-committal about the occurrence of the object of her functional role. Suppose m_j is an employee of the chemical weapons division of a corporation. She desires the failure of the collective act to which she contributes. That is, she does not want the corporation to successfully manufacture and sell chemical weapons to the government. But she also wants to earn a decent income. So upon joining the collective, m_j would have had a protanto motivating reason to agree to perform certain tasks -- but she would also have a protanto motivating reason not to agree to contribute to Σ . This is a coherent position to take, even though the tasks she agrees to perform, ϕ contributes to both Σ and to her private banking account.

The upshot is that if the tasks performed by a member of a collective contribute to both Σ and ψ , and if the member desires that ψ but has no desire that Σ , there is no reason for the member to agree to contribute to Σ . That the member's actions, extensionally described, are consistent with that of a member who has agreed to contribute to a collective act, does not entail that the member agrees to contribute to the collective act. The apparent opacity of agreement prohibits re-describing a member's agreement to perform certain tasks -- which are instrumental to satisfying her private ends -- as an agreement to contribute to a collective act. Yet such individuals are uncontrovertibly members, despite that, it seems, they have not agreed to contribute to a collective act.

The criticism that applied to this principle might also be leveled against the rule-following commitment principle. Why believe that members of a collective actually agree to follow *goal-oriented* rules?

Even if joining a collective requires agreeing to act according to certain rules, it certainly does not require conceptualizing those rules as goal-oriented rules, despite that the rules are indeed goal-oriented. Again, this is because ‘agrees to’ is an intensional predicate. So the fact that acting according to a particular set of rules just is to act according to a particular functional role, does not entail that agreeing to act according to a set of rules just is agreeing to act according to goal-oriented rules. For example, an employer might engineer a set of tasks that are goal-oriented with respect to a collective end, and then instruct an employee to commit these tasks. But the employee need not have any particular intention that has as part of its content a reference to a functional role in general, or to particular *goal-oriented* rules specifically. It seems that to be a member of a collective, it is enough that she agrees to acts in a way that, extensionally described, contributes to a collective act.

Of course, it would be incoherent for an individual to agree to act according to the rules required of her by the terms of employment, while agreeing either *not* to contribute to a collective act, or not to act according to goal-oriented rules. But it is still coherent for an individual to agree to do what is required of her qua member without agreeing -- one way or another -- to act according to *goal-oriented* rules, or to contribute to a collective act.

It might be argued in response that even if “agrees to” has the *semantic* property of intensionality, we can still claim that, for the most part, agreement is closed under known logical entailment. That is, trivial counterexamples aside, if P logically entails Q, and I

know this, then agreeing to P entails intending Q. It is commonly noted that intention is not closed under known causal implication. That is, the fact that P foreseeably causes Q does not mean that intending P entails intending Q. But on my account, following the rules required for membership does not *cause* following goal-oriented rules. Rather, the two are logically related. That is, at least some of the rules required (de dicto) for membership are necessarily goal-oriented. So if joining a collective requires agreeing to act according to rules required for membership, then knowing that these rules are necessarily goal-oriented entails agreeing to goal-oriented rules.

There is at least one problem with this argument, however. Even if “agrees to” is closed under known logical entailment, most individuals do not know that at least some of the rules they agree to abide by are necessarily goal-oriented.

It seems, then, that both the contributory commitment principle *and* the rule commitment principle are false. Given my claim that a commitment to a functional role is required for membership in a collective, many persons who are intuitively members of a collective seem to be illicitly excluded, unless there is some other way that such individuals commit themselves to acting according to a functional role.

4.4.2. DEFENDING CCP AND RCP

In the previous section, I presented arguments against CCP and RCP. According to CCP, membership in a collective entails a commitment to contribute to an appropriate

collective event. According to RCP, membership in a collective entails a commitment to acting according to goal-oriented rules.

The preceding arguments against CCP and RCP are successful only if we ignore the essentially teleological character of participation in collectives. Being part of a collective entails having a purpose in that collective. This purpose, described generally, is to contribute to Σ -- a collective act. Part of what it is to be a member of a collective is to have this purpose. I argue that if an agent voluntarily agrees to be given a purpose in a collective then we can re-describe the content of an agreement regarding the specifics of how she is to fulfill this purpose, in terms of the object of that purpose.

I argue that if a subject 1) agrees to join a collective, 2) agrees to commit the set of tasks ϕ , and 3) knows that ϕ contributes to Σ , and 4) has the purpose of contributing to Σ , then the subject has thereby committed herself to contribute to Σ -- even if she had no particular intention to contribute to Σ . But why believe that, in such a case, we can infer an agreement to contribute to Σ ?

Donald Davidson argued that the teleological nature of intentional acts allows us to re-describe an instrumental act in terms of its intended end, in an opaque context (1978). For example, the intentional act of flipping a light-switch can be redescribed as the act of turning on the light, within an opaque context, so long as the agent's purpose in flipping the light switch was to turn on the light.

For a candidate member of a collective, the teleological character of the role she is adopting in the collective -- i.e., her purpose -- allows us to infer a commitment to contribute to a collective act from an agreement to commit certain tasks. Put in the jargon of my account, the teleological character of the functional role that a candidate member is adopting allows us to infer a commitment to contribute to the object of her functional role, from an agreement to act according to a functional role.

So a member who knows that the tasks she agrees to perform contribute to Σ , cannot agree to have as her purpose a contribution to Σ , and yet coherently and sincerely claim not to have committed herself to contributing to Σ . If such member of a collective sincerely claims that she does not intend to contribute to Σ , then she is in violation of a commitment she has to other members. These claims suggest the following:

Principle of Self-Commitment

1. if j has the purpose of contributing to Σ , and
2. j agrees to commit Φ , and
3. j knows or is in a position to know that Φ is a contribution to Σ , and
4. Φ is in accordance with the goal-oriented rules partly constitutive of j 's functional role,
5. then, j is committed to contributing to Σ .

The upshot is that, on my view, intending to agree to contribute to a collective act is not a necessary criterion for being committed to contributing to a collective act. Membership in a collective does indeed entail a commitment to contribute to a collective event.

But how does an individual come to have contributing to Σ as her purpose? This is, after all, a premise in the argument that an individual can be committed to contributing to a collective act even if that individual does not intend to agree to contribute to that collective act. The easiest way for an individual to have contributing to Σ as her purpose, is for that individual to agree to take on contributing to Σ as her purpose. But when members of a collective -- especially paid employees -- join a collective, their motivating intention might be to enter into an agreement in which an authorized member of the business will provide income in exchange for the performance of certain duties. These intentions might be motivating, as opposed to an intention to be given a purpose in the collective. Suppose j , a prospective employee for a corporation, agrees to perform certain tasks in exchange for a regular income. j refrains from agreeing to contribute to Σ , and from agreeing to take on the purpose of contributing to Σ . She simply agrees to perform certain tasks, and abide by certain rules -- nothing more. Suppose someone who agreed to a purpose which required a contribution to Σ , would commit the same acts as those that j agrees to commit. We cannot infer that j agreed to a purpose which requires a contribution to Σ . Since many individuals who join collectives do not agree to take on a purpose in that collective, it seems that such individuals cannot be said to have committed themselves to contributing to Σ . But I will argue that such individuals are

committed to contributing to Σ , since such individuals have a purpose in that collective -- even if they did not agree to that purpose.

An individual can have a purpose, even if the individual did not conceptualize an agreement to act according to certain rules as an agreement to act according to a purpose. A candidate member of a collective can be committed to acting according to a purpose -- i.e., acting according to goal-oriented rules contributory to a collective act, even if the candidate made no such agreement to do so. I argue that a member has a purpose in virtue of an agreement that she makes to act according to certain rules, which the subject is in a position to recognize as goal-oriented, and in virtue of a disposition of others to recognize her as having a purpose in the collective.

One necessary condition for having a purpose in a collective, is that the subject must agree to act according to rules that she is in a position to recognize as goal-oriented. The individual need not have the occurrent intention of agreeing to act according to goal-oriented rules in order to have a role in the collective, no more than an individual need have an intention to be a public witness in a trial, when an individual knows that she is on the stand swearing to give truthful testimony. In both cases, it is enough that the individual is aware -- or is in a position to be aware -- that she is taking on a role, one which is recognized by others as such. That others in the collective are disposed to believe that an individual has a purpose in the collective, is part of a sufficient condition for that individual to be the subject of a functional role. Another condition is that the candidate member agrees to act according to rules that are in fact goal-oriented. And a

third condition is that the candidate is in a position to recognize the rules as goal-oriented. These three conditions are jointly sufficient for having a purpose in a collective. The conditions for being the subject of a functional role can be put in terms of the following principle, which can be called

The Principle of Public Role Acquisition (PRA):

1. if s agrees to act according to a set of rules, and
2. s knows or is in a position to know that the set of rules is goal-oriented with respect to a collective event, and
3. others who act according to rules with the same end are disposed to regard s as having a purpose in their collective, then,
4. s has a purpose in the collective.

In the following section, I will elucidate how recognition confers a purpose. The first and second criteria ensure that the basis of responsibility for a collective act is grounded in an event that is under the control of the agent -- the voluntary acts consist in an agreement to be the subject of a functional role. (This will be important when discussing individual responsibility for collectively committed acts). The upshot is that, on the Principle of Public Role Acquisition, s need not agree to have a purpose in a collective in order to have such a purpose.

Recall that according to the Principle of Self-Commitment, if a candidate member has the purpose of contributing to Σ , agrees to commit Φ , knows or is in a position to know that Φ is a contribution to Σ , and if Φ is in accordance with the goal-oriented rules partly constitutive of her functional role, then she is committed to contributing to Σ . The Principle of Public Role Acquisition provides conditions for the satisfaction of the first condition of the Principle of Self-Commitment, by attributing to the candidate member the purpose of contributing to Σ . Thus the Principles of Self-Commitment and Public Role Acquisition entail the Contributory Commitment Principle, according to which membership in a collective entails a commitment to contribute to the aims of that collective.

Recall that if we show that both the Rule Commitment Principle and the Contributory Commitment Principle are true, then we have shown that individuals in a collective are committed to acting according to a functional role. Regarding the Rule Commitment Principle, it might be argued that even if having a purpose is a consequence of agreeing to join a collective, the subject is not beholden or otherwise normatively committed to acting according to the purpose partly constitutive of the functional role, because the subject never agreed to do so. It seems that a person is not committed to acting according to a certain purpose merely by having agreed to act in a way that others will perceive as acting for that purpose. But at one remove, this is false. After all, the tasks which she has agreed to perform are co-extensional with conformity to goal-oriented rules. If she does not perform those tasks, then she has violated a commitment -- viz., the commitment to perform those tasks. The commitment to perform certain tasks need not be couched in

terms of conformity to goal-oriented rules in order for a failure to act according to those tasks to count as a violation of a commitment.

Of course, on this view, that it is an individual's *purpose* to commit certain tasks is not her reason to do so. Rather, her reason to do so might simply be that she agreed to commit those tasks. But in doing so she will also be fulfilling her purpose. So it seems the individual could coherently disavow an intention to act according to goal-oriented rules, or to contribute to a collective act, as these might be side-effects of the member's intention to simply perform the tasks she has agreed to perform. She can coherently disavow an intention to act according to goal-oriented rules, or to contribute to a collective act, or, in general, to act according to her purpose in the collective -- all this, in spite of the fact that she knows that performing the tasks she has agreed to perform is extensionally equivalent to acting according to goal-oriented rules, contributing to a collective act, etc. But for such an individual to be committed to acting according to goal-oriented rules, it is not necessary for her to intend to commit herself to acting according to those goal-oriented rules.

So if such an individual sincerely denies an intention to act according to goal-oriented rules -- i.e., to contribute to the object of her functional role -- that is, if she sincerely denies an intention to contribute to Σ , then she is violating a normative commitment. While a sincere denial of her intention to follow through with her commitment is possible, this of course does not change the fact that there is still a (defeasible) moral requirement for her to contribute to the object of her functional role.

I have argued that an individual commits herself to contributing to a particular collective event by acting according to goal-oriented rules, even if she has no such intention to be so committed. I argued that an individual can have a purpose in a collective partly in virtue of the justified attitudes that others have as a result of the agreement that she makes to act according to a set of rules which she is in a position to recognize as goal-oriented.

In sum, I have argued:

1. if j agrees to act according to a set of rules ϕ , and
2. if j knows or is in a position to know that ϕ is goal-oriented with respect to collective act Σ , and
3. if others who act according to rules with the same end are disposed to regard j as having a purpose in their collective,
4. then j has the purpose of contributing to Σ (PPR), and
5. then (from 1-4) j is committed to contributing to Σ , and
6. then (from 1 & 2) j is committed to following rules ϕ that are goal-oriented with respect to Σ , and
7. then (from 5 & 6), j is the subject of a functional role with Σ as its object, and ϕ as its goal-oriented rules.

Call this the ‘The Functional Role Argument’. Note that I have not provided any decisive arguments in favor of the existence of functional roles specifically. At this point, I am

only laying out the view and arguing for its cogency. Arguments in favor of the view will be made in sections 5 and 7.

4.5. FUNCTIONAL ROLE ACQUISITION

According to the Functional Role Argument, if s agrees to act according to a set of rules, and s knows or is in a position to know that the set of rules is goal-oriented with respect to a collective event, and others who act according to rules with the same end are disposed to regard s as having a purpose in their collective, then s has a functional role. But precisely who are these ‘others’ who must be disposed to regard s as having a purpose in their collective? How many such individuals must there be? What if there is disagreement among them regarding whether s has a purpose? Here I elucidate how the relations among individuals in a collective establish functional roles.

4.5.1. CRITERIA FOR MEMBERSHIP

A decision procedure delegating membership helps avoid intractable disagreements over whether a candidate member should be granted membership in a collective. Such a decision procedure would, put roughly, provide the conditions for accepting a candidate as a member of the collective. But what if there is no such decision procedure? What if no member is implicitly or explicitly authorized to accept a candidate as a member? How does a candidate member acquire a functional role in such a case? Call a cooperative collective without any decision procedure delegating membership an ‘informal collective’.

Membership in an informal collective is determined by the attitudes of others in the collective. In a collective without a decision procedure delegating membership, a candidate member is or becomes an actual member only if other members of the collective regard the candidate member as an actual member. That is, the members of a collective who share Σ as the object of their functional roles must regard a candidate member as the bearer of a functional role with Σ as an object in order for the candidate member to be an actual member. This is in keeping with the Principle of Public Role Acquisition.

But such a view raises two questions: 1) what are the attitudes that others must bear in order for an individual to be a member of a collective, and 2) *which* members of a collective must bear these attitudes?

I have answered (1). Put very roughly, an individual i is a member of C iff the other members of C believe that i has a particular kind of purpose. But what about question (2)? That is, *who* in a collective must have the attitudes required for an individual i to be a member of C ? Must *all* members have the appropriate attitudes? This is implausible. Must at least one member have the appropriate attitude? This too is implausible. Perhaps, then, most of the members must have the appropriate attitudes. But this view has been adequately critiqued by Margaret Gilbert (1989, 1994).

There seems to be no percentage of individuals in a collective who must have the relevant attitudes; or, if there is such a percentage, it is not the same for all informal collectives.

And there is no obvious function which takes a collective (or features of a collective) as its input and yields the percentage of individuals who must have the relevant attitudes as its output. Instead, the number of individuals who must have the appropriate attitudes about *i* in order for *i* to be a member depends on the collective -- and even then, the number of individuals will typically be vague. That is to say, there is no number of individuals who must have the appropriate attitudes for *i* to be a member, even though whether *i* is a member depends on the number of individuals who have the appropriate attitudes.

Consider the following analogy. Whether an individual is bald depends on the number of hairs he has. And there are clear instances of being bald and equally clear instances of being not-bald. But being bald does not consist in having less than a specific number of hairs (I am ignoring here Timothy Williamson's Epistemicism). The boundary, as it were, between being bald and not-bald is vague. This means a person is neither bald nor not-bald if that person has a particular number of hairs that falls within the vague boundary separating bald individuals from not-bald individuals. Put philosophically, the maximum number of hairs an individual must have in order to be bald is indeterminate, in that there are fully-specified individuals who fall neither in the extension nor the anti-extension of 'bald'. This is to say that the Law of Excluded Middle does not apply to 'bald' since it is a vague predicate. Even though the maximum number of hairs an individual must have in order to be bald is indeterminate, whether an individual is bald consists in nothing more than how many hairs that individual has.

The identity conditions for being bald are similar to the identity conditions for being a member of an informal collective, in that both are vague. In the same way that the maximum number of hairs required to be bald is indeterminate, the maximum number of members in an informal collective who must have the appropriate attitudes in order for an individual to be a member of that collective is indeterminate. And the same way in which there are unobjectionable instances of being bald and not-bald, there are unobjectionable instances of being a member and a non-member of an informal collective. Even though whether an individual is a member depends on the number of members in the collective who have the appropriate attitudes, there is no *particular* number of individuals who must have the appropriate attitudes. Thus there are individuals who are neither members nor non-members of an informal collective.

So an individual has a functional role with object Σ -- and is thus a member of an informal collective -- only if she is recognized as having such a role by an indeterminate number of others who have Σ as the object of their functional roles. It is partly because the conditions of membership are indeterminate in an informal collective that the members might establish formal or informal decision procedures delegating membership. Otherwise, the membership-status of some individuals might remain indeterminate.

Though the number of individuals with the attitudes required for i to be a member is indeterminate in informal collectives, various features of the collective act affect this indeterminacy, by narrowing, broadening, or shifting the vague boundary between membership and non-membership. For example, in the same way that it is possible for i

to be a member of an informal collective even if there is a portion of the collective that does not regard *i* as a member, it is also possible for *i* to be a member of an informal collective even if a portion of the collective regards *i* as a non-member. The former case is an instance of non-unanimity with respect to attitudes regarding *i*'s status as a member, while the latter case is an instance of disagreement regarding *i*'s status as a member. Typically, disagreement over the membership of *i* counts more heavily against *i*'s membership than non-unanimity with respect to *i*'s membership.

Consider two qualitatively identical collectives -- each has an equal number of members who lack the belief that *i* is a member. However, in collective *D*, the individuals who lack a belief that *i* is a member have the belief that *i* is not a member. Thus they disagree with those who believe that *i* is indeed a member. In collective *N*, the individuals who lack a belief that *i* is a member either have no particular beliefs regarding the membership of *i*, or they suspend judgment. Thus there is non-unanimity regarding the membership of *i*. Even though the number of disagreeing individuals in *D* is the same as the number of non-believing individuals in *N*, and even though *N* and *D* are otherwise the same, it is possible for *i* to satisfy conditions in membership for *N* without satisfying conditions for membership in *D*. Because, all things being equal, disagreement over membership weighs more heavily against membership than non-unanimity weighs against membership, the indeterminate number of individuals who must believe that *i* is a member in order for *i* to be a member is greater in *D* than in *N*.

The belief that *i* is not a member of *C* should be kept distinct not only from withholding or lacking the belief that *i* is a member of *C*, but also from believing that *i ought not to be considered a member of C*. Sometimes it is not possible for a member to coherently regard *i* as actually a non-member. An individual might *commit* herself to the view that *i* is a member of *C*, if the former has a functional role the fulfillment of which requires treating the latter as an actual member. If the member acts according to this functional role, then she treats *i* as an actual member. I do not mean to suggest that it is not possible for one member to treat an individual merely as if that individual were a member, as in play-acting, deception, etc. Intentions, after all, determine whether one member regards an individual as an actual member. But recall that to act according to a functional role is to act according to a purpose. If it is one's purpose to treat certain others as members, then one cannot intentionally act according to this purpose while simultaneously denying that those others are actual members of the collective. In such a case, an intention to act according to one's purpose is incompatible with an intention to regard *i* as a non-member. However, in such a situation, acting according to one's purpose is indeed compatible with believing that *i ought not to be considered a member*.

It might be argued that it is possible for an informal collective to have members unrecognized as members by *anyone* else in the collective. Consider a terrorist group which consists of single-member cells. Each cell acts independently from every other cell, and no member knows any of the other members. It seems, in such a case, that no one in this terrorist group recognizes any one else in the group as a member. And certainly the indeterminate number of individuals required for *i* to be a member cannot

include zero individuals. As an analogy, the indeterminate number of pebbles required to form a heap cannot include zero pebbles (zero pebbles is firmly in the anti-extension of 'heap'). So either I have to deny that the terrorist group constitutes an actual informal collective, or claim that it is possible for *i* to be a member even if *i* is not recognized as such by anyone else in the collective.

But it is perfectly possible, given the example, for *i* to be recognized as a member by any other member, even if not one of them knows who *i* is. Such recognition is possible since each member can refer to any other member by using a complex description rather than by ostention or by name. That is, regarding *i* as a member 'de dicto' counts as regarding *i* as a member. For instance, a terrorist who believes that others who satisfy certain conditions are fellow-members, successfully recognizes as members those who satisfy the requisite conditions -- even if the terrorist does not know *who* satisfies these conditions.

I have described here how an individual might come to have a functional role -- and thus be a member -- in a collective absent formal decision procedures delegating membership. But how is the content of the functional role -- most notably its goal-oriented rules -- established? I turn to this next.

4.5.2. SCANLONIAN EXPECTATIONALISM AND GOAL-ORIENTED RULES

Recall that goal-oriented rules are an integral part of functional roles. In the following section I will argue that one way goal-oriented rules can be established is through implicit agreements. Scanlon's Expectationalist theory of promises (1990, p. 304) can be used to explain how the goal-oriented rules of a functional role may be implicitly established. Scanlon explains the normativity of promising by appealing to what he calls 'Principle F', according to which:

1. if A voluntarily and intentionally leads B to expect that A will do x (unless B consents to A's not doing so), and
2. A knows that B wants to be assured of this, and
3. A acts with the aim of providing this assurance and has good reason to believe that he or she has done so,
4. B knows that A has the intentions and beliefs just described, and
5. A intends for B to know this, and knows that B does know it, and
6. B knows that A has this knowledge and intent,
7. then, in the absence of special justification, A must do x unless B consents to x's not being done.

In an implicit agreement, however, A leads B to expect that A will do x, but A does not provide, nor does B request unequivocal assurance of A's intentions. Though explicit acts of assurance are absent, A nonetheless intentionally leads B to expect that A will do x; i.e., (1) is still true. (2) is epistemically weakened: A reasonably presumes that B wants a presumptive assurance of (1). Instead of providing unequivocal assurance to B that (1) is

true, A intentionally leads B to reasonably believe (1). The remaining premises from (4) onward are epistemically weakened in that knowledge is replaced with reasonable presumption:

1. if m_i voluntarily and intentionally leads m_j to reasonably expect that m_i will follow goal-oriented rules ψ , and
2. m_i reasonably believes that m_j wants to be presumptively assured of (1), and
3. m_i acts with the aim of providing this presumptive assurance and has good reason to believe that he or she has done so, and
4. m_j reasonably believes that m_i has the intentions and beliefs just described, and
5. m_i intends for m_j to reasonably believe (4), and m_i reasonably believes that m_j does in fact reasonably believe (4), and
6. m_j reasonably believes that (5),
7. then, in the absence of special justification, m_i should follow goal-oriented rules ψ , unless m_j agrees otherwise.

As a result, m_i has a prima facie obligation to follow goal-oriented rules ψ . The content of ψ can be ‘fine-tuned’ through iterated applications of this modified version of principle F, by alternating between the negotiating parties.

The resulting obligations can vary in strength and defeasibility depending on the circumstances. And it need not be the case that m_i is related to every other member in this

way -- it is enough if the implicit goal-oriented rules are established with some member, after which other members reference the implicitly established goal-oriented rule.

This is, of course, neither the only method nor a necessary method by which goal-oriented rule can be established -- rather, it is a method by which the content can be *implicitly* established.

5. Metaphysical Conditions Revisited

I have outlined what a cooperative collective consists in, what membership in a cooperative collective consists in, and what acting as part of a cooperative collective consists in. To briefly summarize: a cooperative collective consists in a set of functional roles that share the same object. Membership in a cooperative collective consists in a commitment to be the subject of a functional role constitutive of the pertinent collective. A functional role, in turn, is a three-part relation consisting in a subject, a set of goal-oriented rules, and an object. The subject-particular is the individual who bears the functional role; the subject-description provides the eligibility criteria for a subject-particular. The goal-oriented rules consist of the required actions that the subject is committed to taking, in virtue of being a subject of the functional role. The goal-oriented rules are productive with respect to the functional role's object, which is the collective event that the subject contributes to by acting according to the functional role's goal-oriented rules. An individual acts as part of a cooperative collective just in case that individual acts according to her functional role.

So far I have presented an account of collectives largely by fiat. That is, I have given no decisive reasons to believe that the account is the *right* account of collectives. In what follows, I argue that the account I have presented solves significant problems in the metaphysics of collectives and collective action.

Recall that an account of collective action must explain (or explain away) the apparent normativity associated with collective action. The normativity of individual contributions

to collective action is an integral part of my account. On this account, functional roles explain the basic duties of individuals as members of the collective. I have already argued that individuals implicitly commit themselves to acting according to a functional role. This implicit commitment helps explain why the reactive attitudes of rebuke and resentment are appropriate should an individual in a collective fail to contribute to the collective act (i.e., fail to act according to her functional role). In the following two sections, I will explain how the account addresses the non-extensionality and the sub/supra-compositionality of collective action.

5.1. NON-EXTENSIONAL IDENTITY CONDITIONS

Recall from section 3.2, that some collectives (e.g., corporations, football teams, etc.) can continue to exist through the gradual or simultaneous replacement of all the individuals constituting that collective. It seems, then, that collectives have non-extensional identity conditions. A challenge is to explain -- or explain away -- the apparent non-extensionality of collectives. Here I address this challenge.

Christopher Kutz (2000) addresses this challenge by denying the reducibility of collectives. He argues that the irreducibly holistic nature of collective action explains how a collective can survive changes in membership. According to Kutz, individuals in a collective intend to contribute to an act committed by an irreducibly holistic collective. An irreducibly holistic collective is, in turn, composed partly of individuals, each of whom intends to participate in collective action in order to satisfy what Kutz calls

“overlapping intentions”. Two intentions overlap just in case promoting the satisfaction of one involves bringing about a state of affairs that promotes the satisfaction of another. Each of these intentions aims exclusively at an individual task that contributes to an act committed by an irreducibly holistic collective. Because such a collective consists in part of overlapping intentions, the collective supervenes upon but is not identical to its members; as a result such a collective can survive changes in its membership.

But when Kutz (2000, ch. 3) addresses the concept of irreducible collectivity referenced in the members’ participatory intentions, he is forced to argue that we have reached the limits of analysis -- that some notion of irreducible collectivity is necessary in order to understand collective action. He emphasizes that there is a straightforward sense in which he has provided an individualistic analysis of irreducibly holistic collectives; such a collective is nothing more than individuals with certain attitudes regarding irreducibly holistic collectives. But relegating the reference of irreducibly holistic collectives to the content of the intentions of individuals constituting such collectives, only pushes back the problem of understanding what it means for a collective to act or intend. We need to know what the individual member is referencing in the direct object of his or her participatory intention; and Kutz provides what he argues to be an unavoidably circular understanding of what the individual is referencing.

I argue, instead, that we ought to eschew the notion of irreducible holism altogether. There are, instead, collectives with non-extensional identity conditions that have among their parts functional roles. We can account for the fact that certain collectives can

survive wholesale changes in their membership by identifying the collective with roles rather than with the members of the collective. Such collectives are genuinely reducible because the relevant attitudes in virtue of which the roles exist need not refer to irreducibly holistic collectives.

On the account I presented, a collective is defined by its functional roles. Since collectives are not identical to any combination of the individuals upon which the collective supervenes, the individual members can be replaced without altering the identity of the collective. Take, for example, the Senate of the government of the United States. This collective persists through changes in its membership. For such a collective, a *de re* reference will yield the actual members who contingently compose the collective at a particular time. Alternatively, the collective can be referenced modally *de dicto*, which abstracts the collective from any specific, contingent set of members. On my account, the Senate, referenced *de dicto*, is a set of functional roles with the same object. And each functional role has an open variable satisfiable by anyone who meets certain eligibility requirements. An 'empty' functional role -- i.e., a functional role with no subject-particular -- is simply a non-instantiated functional role. A collective consisting of non-instantiated functional role is a non-instantiated collective. In any case, these functional roles are what stay constant when all the members of the collective have been replaced. Thus there is no need to posit irreducibly holistic collectives in order to explain their non-extensionality.

5.2. NON-COMPOSITIONALITY OF COLLECTIVES

In the previous section I explain how the account I have developed addresses the problem of the non-extensionality of collectives. Here I explain how the account addresses the non-compositionality of collectives.

Recall from section 3.1 that any theory of collective action must account for two characteristics of the relationship between the actions of those who constitute a collective and the actions of the collective itself. I called one of these characteristics “the sub-compositionality of collective action”, and the other “the supra-compositionality of collective action”. The theory I have presented can account for both.

According to the sub-compositionality of collective action, the fact that most or all members of a collective *C* (and no one else) contribute to and cause a collective act does not entail that the collective *C* committed that act. Recall the example of the board of directors of corporation *X* who first vote to reduce dividends, after which they vote to have dinner at a sashimi restaurant. Only the first act of voting had as its subject the corporation, even though the same people performed both acts of voting. Intuitively, this is because the board members, in voting for the reduction in dividends, did so qua members of the board, whereas they vote on where to eat simply as hungry co-workers. On the account I present, to act qua member is to act in accordance with an appropriate functional role.

Each member of the board occupies a functional role in Corporation X. For any arbitrarily chosen board member, the activity of voting on where to eat food can indeed be part of his or her functional role, insofar as it is subsidiary to other duties more explicitly delineated as the object of the functional role. For example, voting where to take a prospective client to eat can be subsidiary to the object of the functional role associated with being a board member. But voting where to eat after work does not fall under the auspices of the functional role in which being a board member consists. Thus a member of a board switches roles between voting to reduce dividends, and voting on where to eat. It is in virtue of having switched roles that the first act of voting is attributable to the corporation while the second act of voting is attributable to a group of co-workers.

Functional roles can also account for the supra-compositionality of collective action. A collective acts just in case an individual bearing a functional role that partly composes the collective acts according to her functional role. The number of persons who act has no bearing on whether a collective acts. If a president or CEO acts alone by exercising an executive privilege, then this act can be attributed to the government or the corporation *tout court*. If a janitor employed by a corporation acts according to his functional role by emptying the trashcans in a hallway, then this act can be attributed to the corporation *tout court*.

This solution to the non-compositionality of collective action also explains how it is possible for a member of a collective to contribute to a collective act without acting qua

member of that collective. For a collective to act, it is not enough that an individual contributes to the object of a functional role -- the individual must so contribute according to a functional role of which she is a subject. If an individual's functional role explicitly or implicitly prohibits a class of acts, and the individual violates that prohibition, then she likely no longer acts in accordance with her functional role, and thus no longer acts qua member of the collective -- even if her act contributes to the object of her functional role, and therefore to a collective act. Put differently, if an act violates the goal-oriented rules partly constitutive of a functional role then the act likely falls outside the scope of that functional role. For instance, suppose a member of the diplomatic corps decides to advance her country's interests through means that violate the rules of her functional role. Though the act is ostensibly committed under the auspices of her functional role as a diplomat, the act is not in accordance with her functional role. If the State Department disclaims the contributions made by the diplomat of the diplomatic corps, then the act cannot be attributed to the State Department (though the State Department would still bear vicarious responsibility for the diplomat's actions).

It is possible, however, for an act that violates the goal-oriented rules partly constitutive of a functional role -- that is, an illicit act -- to fall within the scope of that functional role. That is, committing an illicit act does not necessarily entail acting outside of a functional role. Certain illicit acts fall within the scope of a functional role. Typically there are norms within the collective that govern responses to particular kinds of rule-violations. Put in terms of the account I have developed, a *second-order* functional role provides rules that govern responses to particular kinds of rule-violations. An illicit act can fall

within the scope of the actor's functional role if the meta-rules that govern the violation mandate it as such. For example, if a mid-fielder in a soccer match uses her hands to score a goal with the soccer ball, then she has violated goal-oriented rules (aptly named, in this case). But this does not mean that by using her hands she has acted outside the scope of her functional role and thus ceased acting as a member of her team. This is because there are *meta-rules governing this type of illicit act, according to which the* illicit act should be treated as a punishable offense falling within the scope of the actor's functional role.

This is not to say that to impose punitive or corrective measures for illicitly committed acts is itself to treat the illicit act as falling within the scope of the actor's functional role. A meta-rule imposing punitive or corrective measures might or might not *in addition* treat the illicit act as falling within the scope of the functional role. For example, there are meta-rules in the military governing treatment of soldiers who intentionally use their position in the military to aid the enemy -- according to such a rule, the soldiers have committed treason. According to the norm governing the treatment of soldiers who have committed the rule-violation known as 'treason', the soldier has both violated goal-oriented rules constitutive of his functional role *and* has acted outside the scope of his functional role. This is in contrast to the previous example in which, according to the norms governing the soccer-player's illicit act, the players only violates goal-oriented rules and has not acted outside the scope of her functional role. The fact that in the previous example the illicit act falls within the functional role's scope, while in the latter example it falls outside the functional role's scope, is not an arbitrary feature of the

respective meta-rules governing the illicit acts. Perhaps it is because the treasonous soldier's violation is, in some respect, more egregious that it is treated as falling outside the scope of her functional role. But whatever it is that explains why we have opted for norms that treat particular illicit acts as falling inside or outside the scope of the actor's functional role, it is ultimately the *treatment* of an illicit act that determines whether it falls within the scope of the actor's functional role.

If a norm mandating treatment of an illicit act is ignored, then the actual treatment of the illicit act by those who determine the actor's functional role determine whether the illicit act falls within the scope of the functional role. For example, suppose the President of the United States orders the torture of US citizens without the explicit consent of Congress. However, Congress "looks the other way", as it were, thereby allowing the President to continue committing the illicit acts. Though the President in such a case violates the goal-oriented rules associated with his functional role, it is not the case that he has ceased acting as President. His illicit act falls within the scope of his functional role as President because Congress treats the act as such, in spite of the norms that require treating the President's actions as falling outside the scope of his functional role.

I have not provided necessary and sufficient conditions that determine whether an illicit act falls within the scope of a functional role. But my goal in this section is only to point out that, for a collective to act, it is not enough that an individual contributes to the object of a functional role -- the individual must so contribute according to a functional role of

which she is a subject-particular. This condition might not be sufficient, but it is necessary.

6. Normative Conditions for an Account of Responsibility for Collective Acts

So far, I have argued in favor of my account of collective action by arguing that it solves certain problems in the metaphysics of collective action. Now I will argue that an account of individual responsibility that dovetails with the account of collective action that I have developed, can solve otherwise intractable normative problems in the ethics of collective action. In section seven, I will develop an account of individual responsibility for collective action that dovetails with the metaphysical account of collective action developed in previous sections. I will then show that this account of individual responsibility for collective action can meet the challenges I present in sections 6.1 and 6.2.

I will begin by presenting two conditions, broadly construed, that any account of responsibility for collective action must satisfy. An account of responsibility for collective action that violates either of these conditions must provide a justification for doing so.

First, any account of responsibility for collective action must be compatible with what has been called the “Control Principle”, according to which an individual is responsible for an event only if that individual had relevant control over the occurrence of that event. I discuss this principle in section 6.3.

Second, an account of responsibility cannot result in the disappearance of responsibility for collectively committed wrongs. That is, if a cooperative collective commits a

wrongful act, then there must be agents responsible for that act, even if only partially responsible. This condition can be called “The Responsibility Principle”. I discuss this principle in section 6.2.

Satisfying these two simple conditions is difficult since, in certain cases I will describe, they pull in different directions; satisfying one principle seems to come at the cost of satisfying the other. In what follows, I show how and why this is the case. After elucidating the Control and Responsibility principles and providing reasons for imposing them as conditions, I will show how various existing accounts fail to satisfy them. I will then (attempt to) solve the dilemma that the pair of principles pose by developing an account of responsibility for collective action that builds upon the account of collective action developed in previous sections.

6.2. THE RESPONSIBILITY PRINCIPLE

To understand what motivates imposing the Responsibility Principle as a condition for an account of responsibility for collectively committed harms, it is first necessary to review why such an account is needed in the first place. Why develop an account of responsibility for collectively committed action at all? At one remove, an account of responsibility for actions committed by individuals cannot be applied straightforwardly to actions committed by collectives since collectives are not individuals. But if we have a reductive account of collective action in which putatively collective acts can be understood in terms of the actions of individuals, what remaining reason is there to

develop an account of responsibility for collective action? It seems that a reductive metaphysical account of collective action suffices, in that it permits the application of a theory of responsibility for individual action to collectives.

A theory of responsibility for collective action is needed because applying standard accounts of responsibility to members of collectives yields absurd results. More specifically, oftentimes when we apply standard accounts of individual responsibility to contexts in which an individual acts as part of a collective, we are left with the absurd conclusion that *no one* is responsible for the harm committed by a collective. There is, evidently, something special about collective action that necessitates the development of a theory of responsibility continuous with but also distinct from standard accounts of responsibility.

In what follows, I more closely investigate how responsibility can seemingly disappear if we apply standard accounts of responsibility to contexts in which an individual acts as part of a collective.

6.2.1. THE DISAPPEARANCE OF RESPONSIBILITY

The disappearance of responsibility for collectively committed harms is a consequence of morally evaluating (relative to some axiological schema) acts based on the effects they have. This is a classical problem. Suppose Σ , a collectively committed harm, results from the actions committed by individuals p_1 \dots p_n . Suppose p_1 commits $p_1\chi_1$, p_2 commits

$p_1\chi_1 + p_2\chi_2 + \dots + p_n\chi_n$ result in the collectively committed harm Σ . Assume that if p_i does not commit $p_i\chi_i$, then a backup (b_i) will step in, and commit $b_i\chi_i$. So for any p_i , her contribution to the collectively committed harm is overdetermined, in that her contribution will be made by someone else if p_i opts to refrain from contributing. So whether p_i contributes to the collectively committed harm makes no difference to the occurrence of the harm.

For any p_i , even though her contribution to the harm is overdetermined, *someone* must make that particular contribution in order for the harm to occur. For example, suppose that a dozen individuals together push a boulder off a cliff. Doing so requires a dozen individuals; if fewer than a dozen individuals attempt to push the boulder, the collective act will fail. Suppose that each of the dozen contributors has a back-up; if any particular individual ceases to push, a backup contributor will take her place. Suppose further that each of the backups also has a backup, *ad infinitum*. Thus any p_i 's contribution is overdetermined -- not because the collective effort can be achieved with fewer than twelve individuals, but instead because twelve individuals will contribute regardless of whether p_i contributes. If neither p_i contributes, nor her backup, nor her backup's backup, et al., then the collectively committed harm will not occur.

If we evaluate p_i 's contribution based solely on the difference it makes with respect to the occurrence of the collectively committed harm, then p_i does no wrong by providing an overdetermined contribution to the collectively committed harm. Since each of p_1 to p_n provides an overdetermined contribution to Σ , no individual has done wrong by

contributing to the collectively committed harm, assuming we evaluate each individual's contribution based solely on whether the harm counterfactually depends on that individual's contribution. We thus have a 'disappearance' of responsibility, in that no one is responsible for a collectively committed harm caused solely by the actions of agents who are (ex hypothesi) intentionally contributing to the harm.

It might be argued, however, that p_i 's contribution should be morally evaluated based not on the difference it makes with respect to the occurrence of the collectively committed harm, but rather, based on the *actual* effects of her contribution. By hypothesis, if we rigidly designate the contribution that contingently happens to be committed by p_i , then the rigidly designated contribution is necessary for the occurrence of the collectively committed harm. Put in terms of the previous example, the contribution made by the 6th contributor (for instance) is physically necessary in order to achieve the collective harm of pushing the boulder off the cliff. This is so in spite of the fact that the contribution made by whoever the 6th contributor happens to be is overdetermined, in that someone else would have been the 6th contributor if the actual 6th contributor opted out. If a rigidly designated contribution is necessary for the occurrence of a collectively committed harm, then the individual who *actually* provides that rigidly designated contribution is (at least partly) responsible for that harm -- or so it might be argued.

However, even if we evaluate a contribution to a collectively committed harm based on whether it is an actual instantiation of a rigidly designated necessary contribution, responsibility *still* disappears if the contribution is *not* so small that it is unnecessary for

the occurrence of the collectively committed harm. As a result of the fact that each contribution is so small, each of $p_1\chi_1 \dots p_n\chi_n$ is neither necessary nor sufficient for the occurrence of the collectively committed harm. That is, the collectively committed harm will occur regardless of whether p_i commits $p_i\chi_i$. So even if a rigidly designated contribution necessary for the occurrence of a collectively committed harm is wrongful, p_i does no wrong by committing $p_i\chi_i$, since that contribution, whether or not rigidly designated, is not necessary for the occurrence of the collectively committed harm. So the fact that p_i actually contributes to a harm does not serve as a basis of responsibility for the collectively committed harm, if the actual contribution, designated rigidly, is not necessary for that harm. Again we have a disappearance of responsibility, in cases where the collective harm is the result of the conglomeration of many small contributions.

Of course, some might bite the bullet by accepting the view that $p_1 \dots p_n$ are not responsible for the collectively committed harm Σ . But the view that no single contributor bears significant responsibility for the collectively committed harm, on the grounds that her contribution is overdetermining, is very difficult to accept. This view serves as a *reductio* on act-consequentialism.

In what follows, I present the notion of ‘secondary responsibility’, which, if plausible, will allow us to avoid the disappearance of responsibility in the contexts of collective action. Then I will discuss the factors that affect the degree of secondary responsibility that a member of a collective bears. Following this, I show explicitly how secondary responsibility solves the problem of the disappearance of responsibility.

6.2.2. SECONDARY RESPONSIBILITY

The disappearance of responsibility for collectively committed harms can be avoided by attributing to contributors of the harm what can be called ‘secondary responsibility’. An individual bears secondary responsibility just in case she is morally responsible for the actions of another person (or persons). There are primarily three ways an individual *qua* member of a collective can bear secondary moral responsibility for what a collective does. These three forms of secondary responsibility can be called ‘vicarious’, ‘associational’, and ‘contributory’, respectively. In this section I elucidate these forms of secondary responsibility.

Suppose *p* is a person who commits a harm ϕ . In paradigmatic cases of responsibility, if the only party solely responsible for ϕ is *r*, then *r* is identical to *p*. But this is not always the case. Responsibility for the commission of a harm may fall not on the person who committed the act, but rather on a person recognized as morally responsible for the actions of the wrongdoer. If the only party responsible for ϕ is *r*, and if *p* is not identical to *r*, then *r* is *vicariously* responsible for ϕ .

The legal form of vicarious responsibility is sometimes known under common law as the doctrine of *respondeat superior*, applications of which include an employer’s liability for illicit acts committed by employees *qua* employees, or a superior officer’s liability for illicit acts committed by subordinate combatants *qua* combatants. So if *p* commits ϕ ,

and p is appropriately related to r, then r can be solely responsible for phi. Vicarious liability in common law might be thought of as underpinned by vicarious moral responsibility; like vicarious liability, vicarious moral responsibility does not require the morally responsible party to have been able to prevent the wrongful act. (I will say more about what the scope of this modal claim -- i.e., what counts as being 'able to prevent' -- in section 6.2.2).

In what can be called 'associational responsibility', all members of a collective are responsible for a wrongdoing committed wholly by a proper subset of the members of that collective. Suppose event χ is committed by the collective $G\chi$. However, not all members of $G\chi$ contributed to χ . A member of $G\chi$ who did not contribute to χ (this includes inchoate forms of contribution) bears associational responsibility for χ just in case mere membership in $G\chi$ is enough to make that individual responsible for χ .

Of course, even if mere membership can be the basis of responsibility for a wrongdoing committed by a fellow member, *the degree* of responsibility still needs to be determined. Members can share some responsibility for the wrongdoing without being *as* responsible as the actual wrongdoers. Or members can be fully responsible, in that they are as responsible as the actual wrongdoers. Either way, associational responsibility leaves open the degree of responsibility which the individuals who did not commit the act bear. But whatever the degree of responsibility, bearing associational responsibility settles a *basis* for responsibility, viz., membership in the collective that committed the harmful act.

Closely related to associational responsibility is what can be called ‘contributory responsibility’. Suppose event χ is committed jointly by a group of persons. Call this group “ G^χ ”. Each member of G^χ contributes to χ , and no member commits χ alone. In cases of contributory responsibility, as with associational responsibility, every member of G^χ is responsible for χ . But contributory responsibility departs from associational responsibility in two ways. An individual bears contributory responsibility for χ just in case she both is a member of G^χ *and* contributes to χ . An individual bears *associational* responsibility, on the other hand, if she bears responsibility wholly in virtue of sharing membership with those who contribute to χ . A basis for bearing contributory responsibility is the act of contributing qua member to a collectively committed harm.

For example, suppose a mob boss hires a thug to kill the owner of a local jewelry shop. Even though the thug is the one who pulls the trigger, the mob boss shares with the thug full responsibility for the murder -- that is, both are fully responsible for the murder. Regardless of how their respective contributions are analyzed, it is clear enough that the thug performed the *actus reus*. Though the mob boss was not physically engaged in the act of murder, he provided a significant contribution. Thus both he and the thug together are responsible for the murder. Neither the mob boss nor the thug is responsible merely by association, since each (intentionally) contributed to the wrongful act. As a result, even though the thug performed the wrongful act, they both bear contributory responsibility. Alternatively, if other members of the mob who were not involved in the murder are nonetheless responsible on the grounds that they are part of an organization

committed to this kind of violent activity, then these other members can bear, at most, associational responsibility, since they did not contribute to the murder.

Note, however, that a member of $G\chi$ who bears contributory responsibility for χ does not necessarily bear the same degree of responsibility as every other member of $G\chi$. The claim that a member bears contributory responsibility for χ is a claim about the *basis* of her responsibility for χ .

Recall that my use of ‘responsibility’ is intentionally ambiguous between a reference to blame and moral liability. Thus each of vicarious, contributory, and associational responsibility actually consists of two kinds of responsibility. As with primary liability, secondary moral liability includes liability to compensation, preventive injury, and punishment. Many theorists believe that liability entails culpability. I make no such assumption. It is useful to deny the entailment relation here because secondary culpability is more controversial than secondary liability. Secondary blame has largely been rejected by others as a possibility, for reasons that I will explore and refute in later sections. I will argue that there are compelling reasons, both theoretical and intuitive, to believe that, often, when a collective commits a harm, its members can bear not only secondary liability, but secondary blame as well for that harm.

In what follows, I present some challenges that secondary responsibility generates. I will then argue that these challenges are met by the account of the metaphysics of collective action that I have developed in previous sections.

6.3. THE CONTROL PRINCIPLES: WEAK AND STRONG

A major obstacle for any account of collective responsibility that incorporates secondary responsibility is satisfying -- or providing an error theory for -- what Doug Husak, calls the “Control Principle”, and what I will call the “Strong Control Principle”.

SCP: An individual is responsible for an event χ only if that individual had control over the occurrence of χ .

As a consequence of SCP, if I do not cause χ , then I am not accountable for χ . This is an intuitively plausible principle -- it can be seen as the contrapositive of the principle that ought implies can, viz., cannot implies not-ought.

In what follows, I show how SCP is apparently incompatible with secondary responsibility. I will then argue that SCP is mistaken, after which I will present an alternative to SCP compatible with secondary responsibility.

6.3.1. INCOMPATIBILITY BETWEEN SCP AND SECONDARY RESPONSIBILITY

I contend that secondary responsibility can solve the problem of the disappearance of responsibility -- but to do so, I must first answer challenges to secondary responsibility raised by the SCP. In this section, I show precisely why SCP is deeply at odds with

secondary responsibility. In the next section, I address these challenges, by offering an alternative to SCP which is not only compatible with secondary responsibility, but is also more plausible in its own right.

Suppose p bears contributory responsibility for an event, χ , committed by a collective, G , of which p is a part. Since the responsibility is contributory, we know that p as well as at least one other person contributed to χ , and that χ is the result of these contributions. If χ would not have occurred had p refrained from contributing, then attributing full responsibility for χ to p is consistent with SCP; this is because p could have prevented χ . But suppose the occurrence of χ was overdetermined in that it would have occurred regardless of whether p contributed. Suppose that, moreover, p could not have prevented χ . Can p bear responsibility for the occurrence of χ even though p could make no difference to whether χ occurred?

Suppose we claim that p bears less than full responsibility for χ , since, though p 's contribution could make made no difference to whether χ occurred, his contribution was still part of χ . To claim that p bears a share of the responsibility for χ in virtue of the fact that p contributed to χ as a member of G , is consistent with the claim that p bears contributory responsibility for χ . But since p had no control over whether χ occurred, attributing to p *any* responsibility at all for χ violates SCP.

Of course, while p has no control over the occurrence of χ , he does indeed have control over the particular difference that *his* contribution makes to χ . After all, the world in

which p contributes to χ is distinct from the world in which p does not contribute to χ . Thus we have two distinct events -- χ_1 and χ_2 -- where these events are different possible versions of χ . p has control over which possible version is actualized. The former is a version in which p contributes to the occurrence of χ , and the latter is a version in which p does not contribute to the occurrence of χ . Since p has a choice whether to bring about χ_1 or χ_2 , it is consistent with SCP to claim that the choice she makes is a basis of responsibility for either χ_1 or χ_2 . But this does not mean that p is responsible for χ -- rather, p is responsible merely for the *difference* between χ_1 and χ_2 . That is, p is responsible for the difference that her contribution makes to what G does. This is because p only has control over what *version* of χ occurs. To hold p responsible for the fact that χ will occur *simpliciter*, on the basis that p has control over which version of χ occurs, is to violate SCP, because p has no control over whether χ will occur.

But can we not hold p responsible for the difference he makes -- i.e., the difference between χ_1 and χ_2 ? Which event occurs is, after all, under his control. In such a case, it seems that p would bear partial responsibility for what G does. And attributing this responsibility would not violate SCP. But *ex hypothesi* there is no moral difference between χ_1 and χ_2 from an impartial standpoint -- the two events are, axiologically, just as bad. Since p is responsible, at best, for the difference between χ_1 and χ_2 , and no axiological difference exists between the two, p 's contribution is not a basis of any responsibility for χ . If, contrary to this, we claim that p is partially responsible for what G does because p contributed to χ , then p bears contributory responsibility for an event over which he had no control over, which violates SCP.

A similar problem arises for attributions of associational responsibility in certain circumstances. Suppose p does not contribute at all to χ . However, as a member of G , she shares membership with those who contributed to χ -- or put differently, she is a member of the collective that committed χ . If p bears any responsibility for χ , where the basis of this responsibility lies in the fact that she is a member of the collective that committed χ , then she bears associational responsibility for χ . If, in addition, she does not contribute to χ , and could not have prevented the occurrence of χ , then the claim that she bears associational responsibility for χ violates SCP. One might argue that, in such a case, SCP yields the correct result: that p bears no responsibility for χ . But the purpose of this example is not to show that SCP yields the incorrect result, but rather that the result SCP yields is incompatible with the result that associational responsibility yields. Ultimately, though, I will argue that p does indeed bear associational responsibility in such a case.

It is clear, then, how the Responsibility Principle and the Control Principle can pull in different directions in the context of collective action. To satisfy the Responsibility Principle, the members of a collective must be responsible for the actions committed by the collective of which they are a part. But these individual members often do not have control over the occurrence of the actions committed by the collective of which they are a part. Thus, to attribute such responsibility to them violates SCP.

The prima facie plausibility of SCP presents a challenge for supporters of secondary responsibility. Any theory of collective action that countenances secondary responsibility

must justify - or explain away - the apparent incompatibility between secondary responsibility and SCP. I turn to this next.

6.3.2. THE WEAK CONTROL PRINCIPLE

I argue that, in spite of its prima facie plausibility, SCP is mistaken. The reason why SCP appears true is that it often generates the right results; situations in which an individual does not have control over χ do indeed tend to be situations in which that individual is not morally responsible for χ . But when SCP yields these right results, it does so for the wrong reasons. The reason why can be best understood by comparing SCP with a closely related principle that, I argue, more accurately captures the intuition that individuals can only be responsible for what they can control. This principle can be called ‘Weak Control Principle’.

WCP: An agent is responsible for χ only if the basis of that responsibility lies in an event under the control of that agent.

Christopher Kutz helpfully notes the importance of distinguishing the basis of responsibility from the object of responsibility (2007). The *basis* is that *in virtue of which* an agent is responsible. The *object* is that *for which* the agent is responsible. When we claim that an agent committed a wrong, we might wonder 1) what the wrongful act was, and 2) why the agent is morally responsible for that act. These are two substantively

distinct issues. The first asks for the object of responsibility, the second asks for the basis of responsibility.

Most of the time, the basis of responsibility for some harm ϕ consists of an act as well as the intentions and motivating reasons associated with committing that act. The object of responsibility, in such a case, is the harm ψ . Donald Davidson famously noted that we often describe actions in terms of their intended effects -- e.g., the act of flipping a switch can be appropriately redescribed as an act of turning on the light. As a result, we often describe the basis of responsibility in terms of its object, when the object is an intended outcome of the basis. This practice, though perfectly appropriate, can lead us to mistakenly conflate the basis of responsibility with the object of responsibility.

For example, suppose an individual knows that if he aims and shoots at a particular target, the target is likely to be seriously injured or killed. And suppose the individual intends to kill the target in precisely such a manner, which he does. We typically use the term 'murder' to refer to (roughly) the intentional and unjustified killing of a person. But notice that 'murder' refers to both the basis and the object of responsibility. The object of responsibility is merely the victim's death. And the basis of responsibility consists in acting on an intention to kill the victim. In paradigm cases, a murder requires the right sorts of causal connections between the basis and the object of responsibility.

(Philosophical problems arise when the causal connection between the basis and the object of responsibility is strained. To use a well-worn example, suppose an individual shoots but misses her intended victim by a hair's breadth -- however, the sound of the

gunshot startles a herd of wild buffalo which trample the victim to death. Whether this counts as an intentional killing of the victim depends on whether the unexpected causal connection between the basis and object of responsibility is of the right sort).

Because it is clear that the basis of responsibility must lie within an event under the control of the agent, and because this event is often conflated with the object of responsibility, it is easy to mistakenly think that responsibility requires control over the object of responsibility as well. There is, however, no good reason to believe this as long as we maintain that the basis of responsibility is an event under the control of the agent. WCP states, in effect, that the control principle attaches to the *basis* of responsibility -- not its object. That is, the basis of responsibility -- not the object -- must lie in an event under the control of the responsible agent.

Thus WCP preserves the intuition that an individual must have voluntarily done something in order to be responsible for an event. WCP preserves this intuition without incompatibility with secondary responsibility.

To see this, take, for example, instances of vicarious responsibility. Suppose an employer or a military officer is vicariously responsible for an individual under his or her charge. Suppose the subordinate commits a wrongful act qua employee or qua combatant, for which the employer or commanding officer is responsible. Furthermore, suppose that the vicariously responsible party could not have prevented the subordinate from committing the wrongful act without having agreed to be responsible for the subordinate in the first

place. Though the employer or the officer might not have had control over the occurrence of the event for which they are accountable, they did nonetheless voluntarily commit an act in virtue of which they are responsible for what the employee or the combatant does. For the employer, this voluntary act consists in hiring the employee. For the officer, this voluntary act consists in agreeing to serve as a CO in the military. The basis of responsibility consists in a voluntary act; thus the intuition motivating the Strong Control Principle is preserved without conflating the object of responsibility with the basis of responsibility.

Of course, there are interesting questions regarding what makes an act a basis of responsibility. This is a very deep question; I will not attempt to answer it here. But in the theory I present, I will argue that one particular kind of act can serve as a basis of responsibility for a distinct event. This act is that of joining a collective.

In the next few sections, I will argue that associational and contributory responsibilities are also compatible with the Weak Control Principle. An individual who voluntarily joins a collective that subsequently commits a harm, can bear significant responsibility for that harm even if the individual made only a very small contribution to the collectively committed harm -- or even if the individual made no contribution at all. While the object of the individual member's responsibility is the collectively committed harm, the basis of her responsibility is not her contribution to the harm. Rather, the basis of responsibility is her voluntary decision to join the collective in the first place.

7. An Account of Responsibility for Collective Action

In this section I develop an account of individual responsibility for collective action. This account will make heavy use of the notion of secondary responsibility. I will defend the account I develop chiefly by arguing that it solves the problem of the disappearance of responsibility.

Recall that on the account of collectives that I developed earlier, an individual is a member of a collective just in case that individual has a functional role. And a member has a functional role, just in case she implicitly or explicitly commits herself to acting according to certain goal-oriented rules. I will argue that the basis of responsibility for a collective act is the individual member's decision to become part of the collective -- i.e., to take on a functional role. By agreeing to act according to a functional role, the individual bears responsibility for what the collective does, over and above the responsibility incurred by the difference her contribution makes. And while the individual might not have control over the occurrence of an over-determined collective act, the individual does indeed have control over whether she joins a collective by committing herself to filling a functional role. And this decision is the basis of responsibility for the collective act. Thus the account of individual responsibility that I have presented is compatible with the Weak control principle.

I argue that voluntarily taking on a functional role serves as a basis for responsibility, the object of which is the collectively committed act to which the individual contributes by acting according to goal-oriented rules. Merely in virtue of being a member of a

collective that commits a harm, that member can be partially responsible for that harm even if the member did not contribute to it. Put differently, a member of a collective bears *associational* responsibility for the acts committed by the collective of which she is a part. But why believe that voluntarily joining a collective -- i.e., voluntarily taking on a functional role -- is a basis for associational responsibility?

By the end of section 7, I will argue that joining a collective by filling a functional role constitutes an implicit agreement which serves as a basis for secondary responsibility, in spite of the fact that the object of responsibility is not causally or counterfactually dependent on the member's implicit agreement

7.1. CONTRIBUTORY RESPONSIBILITY

To show that that joining a collective serves as a basis for secondary responsibility, I will first argue that agreements serve as a basis for secondary responsibility, even if the object of secondary responsibility does not causally or counterfactually depend on the agreement.

I begin by arguing that sometimes one person can be responsible for the acts of another if the latter acts at the behest of the former. Suppose x wishes that ϕ . However, x is not able to bring about ϕ . Fortunately, x has an acquaintance, y , who is in a position to commit ϕ . Suppose x asks y to commit ϕ ; y consequently agrees to do so. Suppose that y commits ϕ either because y has a self-interested instrumental or ultimate desire that ϕ or because y has a desire for the satisfaction of x 's desire that ϕ . It is possible for y to

have either of these motivations and still coherently form an agreement with x, in which y promises to commit phi. Suppose that phi is a harmful act. Each of x and y is fully responsible for the occurrence of phi. More to the point: x is responsible for phi even though he did not commit phi. Why is this so?

Clearly, x's desire that phi is not enough to make x responsible for phi. Perhaps the fact that y would not have done phi if x hadn't asked her is what makes x responsible for what y does. This counterfactual dependence can be a basis of x's responsibility for phi, especially since x intends this dependence. But it is, of course, easy to think of counterexamples in which x is intuitively responsible for what y does, even if y's act does not depend counterfactually on x's decision to form an agreement with y. For example, we might suppose that y had more than one decisive reason to commit phi. Suppose Walter wishes to murder his rival; Walter, however, is in a wheelchair, so he is unable to carry out his wish. Walter consequently asks for help from his friend, Frank. Frank agrees to kill Walter's rival. Unbeknownst to Walter, Frank had planned on killing Walter's rival anyway, because the rival had insulted Frank. Frank kills the rival. Each of Frank and Walter is fully responsible for the murder even though the same murder would have occurred the agreement between them. Clearly, if Walter is responsible, it is not because of any counterfactual dependence between his request for help from Frank and the murder.

Perhaps it is because x *actually* caused y to commit phi by forming an agreement with y, that x is responsible for phi. In this case, a causal dependence, rather than a

counterfactual dependence, is what grounds x's responsibility for phi. But consider again the previous example. Suppose that Walter's agreement did not cause Frank to commit the murder. They both meet in secret, and Frank agrees to help Walter kill the target the following week. When the time arrives, Frank forgets about the agreement altogether -- perhaps he is drunk, or perhaps he is full of rage against the target. In any case, Frank kills the target, but not because he agreed to do so. His agreement to do so plays no causal role at all in killing the target.

Recall the question I am attempting to answer -- why is it that x is responsible for what y does when x and y form an agreement in which y agrees to commit ϕ for x? I originally suggested and then quickly rejected the explanation that x is responsible in virtue of a counterfactual dependence between x's agreement with y and what y does. I then suggested that perhaps x is responsible for what y does because x *actually* caused y to commit phi. In response, I turned to an example in which x and y form an agreement in which y agrees to commit phi; though y commits phi, he does this not because of the agreement he formed with x. Is x, in such an example, responsible for what y does? I believe x is indeed responsible. But even if the intuitions of others run in my favor, such intuitions, I suspect, are capricious. Because the relevant intuitions are unreliable, I will not rely on them solely. So is it the case that x is responsible for what y does when x and y form an agreement in which y promises to commit ϕ because x *caused* y to commit phi?

Here is another example that might help show that the agreement between Walter and Frank is indeed a basis of Walter's responsibility, even if the agreement is not what causes Frank to commit the murder. Suppose Walter and Frank do not form an agreement in which Frank promises to kill Walter's rival. And suppose that Walter knows that Frank desires the death of Walter's rival. So Walter points out various reasons for Frank to commit the murder, and he tries to argue against or mitigate the influence of reasons not to commit the murder. Frank subsequently kills the target, but Walter's speech did not causally influence Frank. The speech does not even influence the probability (propensity) that Walter would kill the target. In this case, because Walter and Frank formed no agreement between them, and because Walter's attempt to influence Frank had no effect, Walter is not responsible for Frank's act. Here I hope to invoke the intuition that, supposing Frank would commit the murder regardless of what Walter does, Walter is more responsible for the murder if he tried to bring it about by forming an agreement with Frank, and less responsible if he tried to bring it about merely by attempting to influence Frank. In short, it seems to me that a deviant causal chain is enough to insulate Walter from responsibility for the murder, but *not* if Walter and Frank form an agreement. To further support this intuition, it seems that if Walter agrees to pay Frank to kill the victim, then Walter would still owe the payment even if Frank kills the victim after having forgotten about the promise of payment.²

Nonetheless, I again suspect that intuitions in this example are not decisive. But this is adequate. I need not decisively show that if x and y form an agreement in which y

² I owe this example to Jeff McMahan.

promises to commit ϕ for x, a basis of x's responsibility for what y does lies in the agreement between them, irrespective of whether the agreement was the cause of y's subsequent act. Rather, I need only show that this is a plausible view -- one which will ultimately (I hope) be justified by the ability of my account to solve certain normative problems in the theory of collective action.

Put more completely, I claim that x is responsible for what y does *in accordance* with their agreement, even if it is not because of the agreement that y acts. The secondary liability that an agreement yields depends not on the effects of the agreement, but instead strictly through the act of forming the agreement.

I have presented examples of asymmetric assistance -- where one individual assists another, but not vice versa -- as examples of cases where one individual is responsible for the acts of another. The claim that x is responsible for what y does in virtue of an agreement between them might have greater purchase if the agreement involved symmetric rather than asymmetric assistance. I turn to such cases next.

Suppose x desires the occurrence of χ , and y desires the occurrence of γ . However, x is unable or unwilling to bring about χ , and y is unable or unwilling to bring about γ . Yet each is able and willing to bring about the other's goal. So x agrees to commit γ for y, and y agrees to commit χ for x. Patricia Highsmith's story "Strangers on a Train" exemplifies such an arrangement. In the novel, Haines has a motivating reason to kill his wife, and Bruno has a motivating reason to kill his own father. While on a train, the

characters meet, and learn of each other's situation. Bruno, as a result, makes the following offer to Haine: Bruno will kill Haines' wife, if Haines kills Bruno's father. Since neither Bruno nor Haines has any apparent motive to commit the murders that Bruno suggests, the police will have no reason to suspect either of them. In such a case, x agrees to commit γ for y, and y agrees to commit χ for x.

Note that χ and γ together is not a salient event. As a result, neither x's act nor y's act counts as a contribution to a jointly desired end. This can be called an instance of 'plural symmetric assistance'. Since x agrees to commit γ for y, and y agrees to commit χ for x, x bears secondary responsibility for what y does, and y bears secondary responsibility what x does.

Contrast this with a case in which x and y both wish that Σ . However, neither can bring about Σ on her own. Suppose χ and γ are acts that promote Σ . Jointly, but not severally, they are sufficient for the occurrence of Σ . Suppose x is in a position to commit χ , and y is in a position to commit γ . As a result, x and y can bring about Σ together; that is, if x commits χ and y commits γ then Σ will result. Suppose x and y agree that x will commit χ and y will commit γ . As a result, they together bring about Σ . I will call a case of symmetric assistance where each party has one and the same goal, 'singular symmetric assistance'. In such a case, I claim that x is responsible for what y does in furtherance of Σ , and y is responsible for what x does in furtherance of Σ , provided that each agreed to act together. That is, x is responsible not only for χ , but for γ as well, and y is responsible

not only for γ , but for χ as well. Thus each of x and y is fully responsible for both χ and γ , and thus Σ .

One method of justifying these conclusions involves arguing that because x would not have acted without y , x shares responsibility for what y does. And because y would not have acted without x , y shares responsibility for what x does. This proposed justification for attributions of secondary responsibility is analogous to the initial justification for secondary responsibility provided for cases of asymmetric assistance. Put generally, if x and y form an agreement, then each is responsible for the (expected) outcomes of that agreement. In the example of asymmetric responsibility, the expected outcome was simply ϕ . In the current example of symmetric responsibility the expected outcome is χ , γ , and Σ . Because of this, it might be argued, each party who entered into the agreement is responsible for its expected outcome because without the agreement neither χ , γ , nor Σ would have occurred.

A counterfactual dependence between the agreement and the occurrence of χ , γ , and Σ can indeed be a basis of x 's and y 's responsibility for those events. But a counterfactual dependence is not the sole basis of responsibility. As in cases of asymmetric assistance, it is not difficult to imagine examples in which x and y are intuitively responsible for χ , γ , and Σ , even though at least some of those events do not depend counterfactually on the agreement between x and y . Suppose x will commit χ regardless of whether she enters into an agreement with y who cannot prevent x from committing χ . Assume that x mistakenly believes that χ is sufficient for the occurrence of Σ . If x and y form an

agreement in which, to achieve Σ , x promises to commit χ and y promises to commit γ it still seems that not only is x responsible for committing χ , but y is responsible for χ as well, even though y could not have prevented χ . So if y is responsible for what x does in furtherance of Σ , the basis of responsibility does not lie solely in a counterfactual dependence between the agreement and χ .

As with the case of asymmetric assistance, it might be argued that a mere causal dependence between the agreement and χ , γ , and Σ is a basis of responsibility for those acts. But again, we can suppose that though x and y form the relevant agreement, and though x subsequently commits χ , the agreement is not what motivates the commission of χ . I claim that in such a case, x is still responsible for γ , and y is still responsible for χ , where a basis of this responsibility is the agreement x and y have made to achieve Σ . This agreement is the basis of responsibility even if the agreement turns out not to be the cause of χ or γ (and thus Σ). Again, each of x and y is responsible for what the other does within the scope of their agreement not simply because their acts are an effect of the agreement, but also because both individuals chose to enter into an agreement in which they have promised to act *for* one another. The object of secondary responsibility that the agreement yields is determined not just by the effects of the agreement, but also through the act of forming the agreement itself.

The agreement that serves as a basis of x's secondary responsibility for what y does (and vice versa), need not be explicit. In section 4.5.3 I explained how Scanlon's theory of promises can be modified to provide an account of implicit agreement. In an implicit

agreement, x intentionally leads y to reasonably expect that x will do χ . Since x has good reason to believe that y would like presumptive assurance that x will do χ , x acts with the aim of providing this presumptive assurance, and has good reason to believe that she has succeeded. And y has good reason to believe that x has good reason to believe that she has successfully acted with the aim of providing y with presumptive assurance that x will do χ . And x intends for y to have good reason to believe that x has good reason to believe that she has successfully acted with the aim of providing y with presumptive assurance that x will do χ . If x has good reason to believe that this intention is successful, and y has good reason to believe both that x has this intention and that x has good reason to believe that this intention is successful, then x and y have formed an implicit agreement in which x has implicitly promised to commit χ .

Note that I have not given necessary and sufficient conditions for bearing secondary responsibility. I only claim that if x and y form an agreement in which y agrees to assist x, this agreement is enough to act as a basis of secondary responsibility in that x is responsible for what y does in furtherance of x's goals.

The types of symmetric assistance I have covered so far, both plural and singular, are still very rudimentary since they involve only two individuals. Consider a slightly more complex case. Suppose set P consists of persons p_1 through p_n . Suppose the members of a subset of P, call it P', desire the occurrence of Σ . No single individual can bring about Σ on her own. Rather, many individuals must contribute in order for Σ to occur. So all members of P' agree with one another to contribute to the satisfaction of Σ . However,

suppose these contributions are not enough to yield Σ . As a result, members of P who are not members of P' offer assistance to members of P' -- specifically, they offer to contribute to Σ . We can assume that members of P offer this assistance, not because they desire that Σ , but because the satisfaction of the desires of P' for Σ is instrumentally beneficial to the rest of P. Members of P' agree to the contributions made by the rest of P. Thus, in this example, members of P' agree with one another as well as with the other members of P to contribute to Σ . In addition, the rest of P agree with one another as well as with members of P' to contribute to Σ .

Though the number of persons is greater in this example than in previous ones, the basic structure is the same: some individuals form an agreement with other individuals, in which the former promise to contribute to a particular goal shared by the latter. Now that I have elucidated how agreements can form the basis of secondary responsibility among groups of individuals, I can make a more general claim:

The Principle of Secondary Responsibility in Collectives (PSR) Each member of a collective bears some secondary responsibility for the actions committed by every other member qua member, in virtue of entering into a singular symmetrical agreement with every other member.

For example, suppose a dozen individuals form an agreement to set off a bomb together. I am one of those individuals. By making this agreement, I have in effect entered into *multiple* agreements -- one agreement with each of the eleven individuals. In each case I have committed myself to helping and to be helped by that individual in furtherance of setting off the bomb. And in each case I will bear some secondary responsibility for the contributions made by that individual. Since the aggregate of these contributions is the event consisting in setting off the bomb, and since I bear some secondary responsibility for the contributions made by every other individual, I will bear some secondary responsibility for setting off the bomb. The basis of this responsibility is my *agreement* to help set off a bomb rather than my actual contribution; the same is true for every other member. Since the basis of the secondary responsibility I bear is my agreement rather than my contribution, I will bear some secondary responsibility for what the others do even if I ultimately make no actual contribution to setting off the bomb. According to the Principle of Secondary Responsibility in Collectives, such agreements -- or at least the commitments that they entail -- are implicitly present in any cooperative collective.

The Principle of Secondary Responsibility in Collectives might seem too strong. In the next subsection, I will clarify and defend the Principle. After this, I will explain how secondary responsibility is not “all or nothing” -- it can come in degrees. The ultimate goal, though, is to show that secondary responsibility prevents the disappearance of individual responsibility in collectively action.

7.2. PRIMA FACIE ARGUMENT AGAINST THE PRINCIPLE OF SECONDARY RESPONSIBILITY

I am developing an account of individual responsibility for collectively committed acts, according to which a member of a collective bears secondary responsibility for the collectively committed act. In the previous section, I presented one basis for bearing secondary responsibility, viz., agreements of a certain sort. More specifically, I suggested that each of asymmetrical agreements, plural symmetrical agreements, and singular symmetrical agreements, forms a basis for secondary responsibility. Here I argue that members of a collective implicitly enter into singular symmetrical agreements, in which the event they have committed themselves to promoting is the shared object of their functional role.

It might seem perfectly possible for an individual to become a member of a collective without entering into symmetrical agreements with each and every other member.

Moreover, if a member does not want to bear secondary responsibility for the actions of those who are contributing to the shared object of her functional role, then, if she is rational, she will have no motivating reason to enter into the relevant singular symmetrical agreements with her fellow members. Or if the member has an active desire *not* to bear secondary responsibility for the contributions made by other members to the object of her functional role, then it seems she has a motivating reason to refrain from entering into singular symmetrical agreements with other members.

For example, suppose *h* is a morally conflicted employee working in the chemical weapons division of a corporation. It is intuitively coherent to claim that, in spite of his

contributions, he lacks a desire for the success of the goals of his employers. It is, moreover, intuitively coherent to claim that, in spite of his contributions, he hopes that his employers fail to meet their goals. He remains an employee, and contributes to the object of his functional role, solely to earn a living wage. Thus h has no motivating reason to enter into an agreement in which he accepts the contributions made by others in the collective toward the object of his functional role. Thus it seems that an agreement with other members of the collective cannot serve as a basis of secondary responsibility for h.

It seems, on my account, that either the Principle of Secondary Responsibility in Collectives is false, or h cannot be a member of the collective consisting of the corporation employing him. But the latter disjunct is absurd -- certainly h is a member of the collective employing him. In fact, there is nothing distinguishing him from other members, other than his moral compunctions and his hopes regarding the failure of the collective of which he is a part. It might seem, then, that the Principle of Secondary Responsibility must be false.

It is also not difficult to envision scenarios in which a member has no reason to enter into a singular symmetrical agreement with other members, *even if* she desires that others contribute to the object of her functional role. Suppose k is instrumentally rational, and k desires that Σ . She correctly believes that by contributing to Σ she increases the chances of the occurrence of Σ . Thus she intentionally contributes to Σ . Moreover, since she is rational, she desires that others contribute to Σ as well. But suppose Σ is an illicit

collective act, such as robbing a bank; k would have little reason to enter into an agreement with the other participants, as this would increase her own responsibility for Σ . This is in spite of the fact that k desires that others contribute to Σ . Again, if being a member of a collective entails entering into a singular symmetrical agreement with the other members of the collective, then it might seem that either, in this example, k is not a member, or the Principle of Secondary Responsibility is false.

The upshot is that it seems the Principle of Secondary Responsibility is problematic, because it seems that an individual can voluntarily take on a functional role, and thereby agree to contribute to the object of that functional role without entering into a singular symmetrical agreement with other members of the collective, perhaps because the member does not value or disvalues the object of the contribution, or perhaps because the member simply does not want to be responsible for the object of contribution. So why believe the Principle of Secondary Responsibility?

7.3. DEFENDING THE PRINCIPLE OF SECONDARY RESPONSIBILITY

In defense of the Principle of Secondary Responsibility, I will argue that because all the members of a cooperative collective share the purpose of contributing to the shared object of their functional roles, membership in a collective involves an implicit agreement to provide and accept help from other members. As a result, for any member, her membership is a basis for bearing secondary responsibility for the contributions made

by the other members of the collective. Such a member cannot coherently deny an intention to provide or accept assistance from others.

I have argued in section 4 that membership in a collective entails a (defeasible) commitment to contribute to a functional role's shared object. Here I will argue that these mutual commitments either *just are* implicit symmetrical agreements to contribute to a shared end, or are tantamount to such agreements. That is, I will argue that joining a collective entails agreeing to act according to a functional role, which in turn entails part of a singular symmetrical agreement with other members. An individual cannot be a member of a collective without being committed to other members in the collective, where the commitment is identical or tantamount to the commitment resulting from a singular symmetrical agreement.

I also mentioned in section 4 that members of a collective have a function, in virtue of being members. This feature of membership in a collective will help explain why and how each member implicitly agrees to provide and accept assistance toward the shared object of her functional role.

When an individual pursues a personal project outside the context of a cooperative collective, that individual does not implicitly agree to accept the assistance of others. So why do contributions to the object of one's functional role -- not to mention mere membership in a collective -- entail an implicit agreement with other members of the collective? The fact that individuals in a collective cooperate and act in a coordinated

fashion means that agreeing to join a collective entails agreeing to provide and accept assistance with respect to a shared object -- or so I will argue.

If j and k are members of the same collective, their contributions to Σ are coordinated. This does not necessarily mean that their contributions causally interact. A limiting case of coordination between the contributions made by j and k requires no interaction between them -- but their contributions are still coordinated in that *if* the optimal state of affairs for yielding Σ required interaction between j and k , then their goal-oriented rules would require such interaction. Such a goal-oriented rule is *dispositionally* interactive. It just so happens that yielding Σ does not actually require interaction between j and k . Instead, the state of affairs required for yielding Σ is extensionally equivalent to a case where j and k ignore each other.

Of course, among the goal-oriented rules that do not require interaction with any members, not all of them are dispositionally interactive. But goal-oriented rules that do not require interaction with any members are nonetheless characteristically dispositionally interactive, in that we would expect an explanation as to why a particular functional role would not require its bearer to contribute in the event that contributing required interaction with other members.

Coordinated action, in general, is the means by which a collective event is achieved in most collectives. Most of the time, the various actions that individuals commit qua members of a collective are intelligible only if we recognize that these individuals are

acting in a coordinated manner. To an alien anthropologist lacking the concept of coordinated action, it might seem remarkable that I am able to consistently anticipate how others will act, when deciding how I should act in the context of a cooperative collective. That the actions of individuals in a collective are mutually responsive in this way is, of course, no coincidence; well-conceived functional roles will include goal-oriented rules that require coordinated action, when such action is instrumental to the shared object of the functional roles.

The fact that coordinated action in a (well-functioning) collective is productive with respect to the objects of functional roles, provides a motivating reason for at least some of its members to coordinate action in the first place. To an alien anthropologist who has the concept of coordinated action but who does not have the concept of cooperative collectives, it would be a mystery *why* members of collectives coordinate their actions. It might seem like an elaborate dance, in which the actions of individuals are causally sensitive to one another, but for no apparent normative reason. But it is no mystery; coordination is instrumental to bringing about a collective event.

If an alien anthropologist wishes to know what reasons a member, *k*, has to coordinate her actions with others, we might say that, for example, she has agreed to do so, or that she has an ultimate desire to do so, or that doing so is instrumentally valuable to *k*, or that she values Σ . These reasons invoke *k*'s intentions, motives, and desires, and the role that contributing to Σ plays in their satisfaction. But in addition to these reasons, there is another reason for *k* to contribute to Σ -- it is her *purpose* to do so. All the members of a

collective, whatever their specific tasks, have an over-arching, univocal purpose to contribute to that collective event. This purpose provides a protanto reason for an individual, once she is a member, to contribute to the collective act Σ , where the contribution is properly guided by the goal-oriented rules in which her functional role partly consists.

This general purpose shared among all members of a collective provides a *domain-restricted, agent-neutral* reason to contribute to Σ . The reason is domain-restricted in that the description of the reason includes an essential reference to the collective the members of which the reason is for. But the reason is agent-neutral within that domain, since the reason does not logically restrict its application to any particular individual. So if a member k has a domain-restricted agent-neutral reason to contribute to Σ , then any arbitrarily chosen individual in the collective has a prima facie reason to contribute to k 's contribution to Σ (notwithstanding functional roles that prohibit arbitrary contributions).

This would not be the case if the only reason k has to contribute to Σ were, for example, provided by an agent-relative desire, such as a desire to get paid. That someone *else* contributes to Σ does not satisfy k 's desire to get paid. There is, thus, no reason for k to desire that the other member contributes. But the agent-neutral reason to contribute to Σ does indeed provide a protanto reason to desire that others contribute -- even if k does not respond to this reason.

What does this have to do with symmetrical agreements in a collective? Recall from section 7.1 that when one individual intentionally assists another, and the assisted individual accepts this assistance, the assisted individual bears secondary responsibility for the assistance. The contribution made by one member need not causally promote another's contribution, or vice versa, in order for the former contribution to count as assistance to the latter. That is, even if neither k nor her contribution interacts with j or his contribution, both contributions can still count as acts of assistance to the other. It is enough if both are contributions to Σ made in accordance with functional roles that have Σ as their objects. Since it is k 's purpose to promote Σ , any contribution to Σ helps k . And since it is j 's purpose to promote Σ , any contribution to Σ helps j .

Since all members of a collective have a domain-restricted, agent-neutral reason to contribute to Σ , any contribution made by a member k counts also as assistance to member j -- provided that j *accepts* the contribution *as* assistance.

I am using "acceptance" as a performative term; for p to accept ϕ is to provide a signal inviting others to act in certain ways that might not be permissible had p not provided the signal. It is the signal that permits q to provide assistance. Put differently, since accepting help is performative, whether p accepts help from q does not depend on whether p actually *wants* q to provide assistance. If p signals a permission for q to provide assistance, (and if this signal is not coerced, made under duress, etc.) then p cannot claim that q provided assistance without permission, even if p did not desire q 's assistance.

But this raises a difficulty. I have claimed that by agreeing to be part of a collective, a member accepts assistance from other members in furtherance of the shared objects of their functional roles. This includes contributions made by others that would have been made regardless of whether the member had joined the collective. So at one remove it might sound odd to say that the joining member 'accepts' the assistance of others, since this performative has no bearing on whether the subsequent assistance occurred.

Typically an individual accepts assistance in cases where the assistance would not have been provided had it not been accepted. Though this dependence applies in paradigm examples of accepting assistance, it is not true in all cases. Suppose that p is engaged in activity the accomplishment of which requires assistance from another person. Such a person, q, offers such assistance. If p signal acceptance of his offer, he will provide assistance. If p doesn't signal acceptance, he will not provide assistance. In this case, whether q provides assistance depends on whether p signals that she accepts his assistance. So, provided that p is sincere, q will provide assistance if and only if p accepts his assistance. But suppose, alternatively, that q will provide assistance to p whether or not p accepts his assistance. Moreover, suppose that p know that there is nothing he can do to prevent q from providing assistance. Nevertheless, in this case, p can coherently consent to q's assistance. And this consent has the same normative consequences that it would if q's assistance did in fact depend on q's consent. For example, if p accepts q's assistance, then p cannot subsequently claim that q provided assistance without q's permission -- even if q would have done so had p not accepted the assistance.

The upshot is that it is perfectly coherent to claim that a member accepts assistance from other members in furtherance of a shared end, even if the event in which the assistance consists would have occurred anyway. And to (performatively) *reject* assistance completely, a member of a cooperative collective would have to leave the collective, since membership entails acceptance of assistance from other members. This is not to say, however, that leaving the collective is necessary to morally condemn the contributions made by others. It is possible to performatively accept assistance that one condemns. I will address this possibility shortly.

So far I have only argued in favor of the coherence of the claim that members of a collective accept each other's contributions in virtue of being members. But why believe this claim in the first place?

I claimed that the acts committed by members of a collective are intelligible only if we regard the members as acting with the purpose of bringing about a certain end, where a means to bringing about that end is acting cooperatively; i.e., assisting one another. Acting cooperatively is a means to the satisfaction of the purpose shared by all members. The contributions made by j are implicitly accepted as help by k , in virtue of the fact that both j and k have as their purpose promoting Σ cooperatively. Working with others in order to bring about Σ entails accepting the contributions to Σ made by the individuals with whom one cooperates, since such contributions further the shared purpose of bringing about Σ . Whatever k 's personal reasons for contributing to Σ , it is her purpose to coordinate and cooperate with others in order to bring about Σ . Having voluntarily taken

on the purpose of promoting Σ cooperatively entails implicit acceptance of assistance from other members. For an individual, then, to voluntarily join a collective and yet disclaim any intention to accept or provide assistance belies the claim that the individual is a member of the collective. Such a claim would be disingenuous at best, and incoherent at worst.

Of course, for an individual member j to implicitly accept contributions made by others does not mean that j must, all things considered, morally approve of the contributions made by fellow members in the collective. We need but imagine collectives in which some members disapprove of the contributions made by other members to their shared object. In such a case, it is not the object that the disapproving members disapprove of, but rather the particular means by which some of the other members contribute to that object. That it is possible for a member to morally disapprove of the contributions made by others does not imply that members do not implicitly accept one another's contributions. Both disapproving and accepting a contribution made by another, though coherent, is a normatively questionable position to take, since the moral disapproval of a contribution provides a strong reason not to accept it as assistance. Perhaps this is a reason why an individual member of a collective might be uncomfortable if there are other members whose contributions are morally suspect. This moral disapproval is at odds with the acceptance of their contributions implicit in sharing a common purpose.

In summary, joining a collective entails a general purpose to promote the object of a functional role, Σ . The fact that every member has as her purpose the promotion of Σ ,

allows us to regard any contribution to Σ as an act of assistance to others in the collective, in furtherance of the purpose to promote Σ . That it is one's purpose to coordinate and cooperate in order to promote Σ entails implicit acceptance of the contributions to Σ made by the individuals with whom one cooperates.

I argue that since each member of a collective implicitly agrees to provide assistance to and accept assistance from every other member of the collective, where the object of this assistance is the achievement of Σ , each member is responsible for Σ . This is true, even if no individual committed Σ , and even if no individual contributed substantially to Σ . Since each member severally assists and/or severally assisted by every other member, and since each member bears secondary responsibility for each act of assistance provided by other contributing members, and since these contributions together are equivalent to Σ , each member therefore bears secondary responsibility for Σ . This is not to say, however, that each member is *fully* responsible for Σ . (I turn to this issue in section 7.4).

7.X. THE RE-APPEARANCE OF RESPONSIBILITY

Now that the account of individual responsibility for collectively committed harms has been mostly laid out, I can show how the account prevents the disappearance of individual responsibility in the context of collective action (as described in section 6.2.1).

Recall that if an individual is responsible only for the events under her control, then there are many collectively committed acts for which no individual is responsible. For example, if 1) no individual's contribution is necessary for the occurrence of a collective

act, 2) no single individual contributes substantially to the collective act, and 3) no individual can prevent the occurrence of the collective act, then there is no individual who is responsible for the occurrence of the collective act. At best, each contributing individual is responsible for her own contribution, which is, *ex hypothesi*, not a substantial portion of the collective act.

In agreement with Christopher Kutz, I argued that we should abandon the Strong Control Principle, which limits responsibility to events under the agent's control, in favor of the Weak Control Principle, according to which only the basis of responsibility must be under the agent's control. Under the Weak Control Principle, the object of responsibility can include events which the agent did not cause and could not have prevented, so long as the agent voluntarily committed an act that made her responsible for the event. When one individual is responsible for the actions committed by a second individual, and the basis of this responsibility does not lie in any causal dependence between the first individual and the act for which she is responsible, then the responsibility which the first agent bears can be called "secondary responsibility". Secondary responsibility is not possible given the Strong Control Principle, but is indeed possible given the Weak Control Principle.

I argued that an implicit agreement to help and to be helped in furtherance of a particular end entails some responsibility for the contributions toward that end by the individuals with whom one has made the implicit agreement. The basis for this responsibility is the (implicit) agreement to help and to be helped, rather than a causal relation between the

agreement and the contributions made by those with whom the agreement is made. As a result, this responsibility is secondary. I argued that joining a collective involves an implicit commitment to the other members of the collective -- a commitment to help them and to be helped by them in furtherance of a shared goal (i.e., the ultimate object of the member's functional role). Thus each member of a collective bears some secondary responsibility for the actions committed by every other member qua member, in virtue of entering into a singular symmetrical agreement with every other member.

Secondary responsibility has as its object the actions committed by some other *individual*; but it's possible for a member to bear secondary responsibility for the collective act 'tout court', even though there is no individual who committed this act. If an individual bears full secondary responsibility for the contributions made by each and every member, then the secondary responsibility that the individual bears is tantamount to secondary responsibility for the collective act tout court, since the collective is nothing more than the individual contributions made by each member.

Of course, each contributing member will bear primary responsibility for his or her causal contribution to the collective act. If the contribution is causally insignificant then the member's primary responsibility is likely to be causally insignificant as well, (which is why responsibility for certain collective acts 'disappears' in the first place). Primary responsibility is limited by the effects that the agent has on the collective; secondary responsibility has no such limitation. Each member bears some secondary responsibility regardless of how much -- or whether -- he or she contributes to the collective act. The

members bear secondary responsibility in virtue of the implicit commitment that they have made to one another in virtue of joining the collective.

The upshot is that a member of a collective can bear significant responsibility for a collective act, even if her contribution is causally insignificant. Responsibility for a substantial collectively committed harm consisting wholly of causally insignificant contributions disappears only if we ignore secondary responsibility as well the commitments inherent in joining a collective. Each member composing such a collective, regardless of the extent of her contribution, bears secondary responsibility for the contributions made the other members, where the basis of this responsibility is the act of joining the collective. Hence, responsibility reappears.

Consider the following example. Peter works for Evil Inc., which generates profits by selling cigarettes to minors in under-developed countries. Peter takes the job because it pays well - not because he has any particular desire to contribute to the sale of cigarettes. However, contributing to this goal, by acting according to his functional role, is a means to getting paid. Thus Peter has an instrumental intention to contribute to the aim of selling cigarettes - i.e., the object of his functional role. In virtue of voluntarily taking the job and accepting a functional role which has as its end the sale of cigarettes to minors, Peter has implicitly agreed to provide and accept assistance, as delineated by his functional role, toward the end of selling cigarettes to minors. Peter is thus responsible for what others do, when they help satisfy the object of his role (provided they act in accordance with their respective functional roles). And others in the collective, as a result of accepting

help from Peter, are responsible for his contributions to the satisfaction of the object of their role. Thus Peter is responsible for what the members of Evil Inc. do together, even though his own contribution is marginal.

But it is clearly unfair to claim that Peter is fully responsible for the harms caused by all the members of Evil Inc. The view I am presenting, however, does not have this consequence. I turn to this issue next.

7.4. DEGREE OF SECONDARY RESPONSIBILITY

There are various factors affecting the degree of secondary responsibility that a member bears. For example, though secondary responsibility has as its basis an implicit agreement to join a collective, actually contributing to a collective can affect the degree of secondary responsibility that an individual bears. All things being equal, if a member contributes to a collective act by acting according to her functional role, then she bears greater secondary responsibility. The contribution is not a basis of secondary responsibility; voluntary membership is the basis, and contributing, in general, strengthens this membership.

Another affecting the degree to which a member bears secondary responsibility for the acts committed by others, is the member's intentions. The account of individual responsibility for collective action that I have developed is sensitive to differences in intentions that are typically relevant to responsibility for acts committed outside the

context of cooperative collectives. The intentions an individual has toward the acts she commits qua member, and the attitudes she has toward the collective act to which she contributes, affect the member's primary responsibility for those acts, and secondary responsibility for the acts committed by others. In section 7.4.1 and 7.4.2, I explain how these intentions affect the degree of secondary responsibility for the acts committed by others in a collective.

The way in which one influences the contributions made by others in a collective can be relevant to the degree of secondary responsibility that the former bears for the actions committed by the latter. For example, suppose an individual forms a collective with others by coercing them into contributing to a collective act. All things being equal, the degree of secondary responsibility that the individual bears for the actions committed by the others in the collective will, in such a case, be higher than it would have been had he merely requested volunteers. Various forms of influence -- coercion, deception, enticement, encouragement, etc., can affect to varying degrees the secondary responsibility that an individual bears for the actions committed by those whom he influences. Likewise, the victims of certain kinds of influence might bear *less* secondary responsibility for the actions committed by others similarly victimized in the collective. These forms of influence might occur within a collective independently of any implicit or explicit agreements that form the basis of secondary responsibility. But when coercion, deception, etc. occur within the context of the implicit agreement associated with joining a collective, then the affects that these forms of influence have on secondary responsibility is more widespread. I focus on this issue in section 7.4.3.

Various forms of influence can affect not only the degree of secondary responsibility, but also the degree of primary responsibility that an individual bears. Suppose a leader in a collective coerces subordinates into committing certain acts. The leader will bear significant secondary responsibility for the actions committed by the others, in virtue of the fact that, because of the leader's intended coercive influence, those actions were not freely agreed to by those who committed them. The leader also will bear significant *primary* responsibility for the difference that she makes in the world -- where the actions committed by the subordinates are among those differences. I will discuss partial primary responsibility in section 7.5.

There are two dimensions along which secondary responsibility can be measured, when measuring individual responsibility for collective acts. The *radius* of secondary responsibility concerns the absolute number or percentage of individuals in the collective for whose contributions a given member bears secondary responsibility. I have argued that a member bears secondary responsibility for the contributions made by *all* other members of the collective. One might argue, however, that it is possible for an individual to have a functional role that limits the radius of secondary responsibility -- or that it is possible for a member to have a functional role that limits whether others in the collective bear secondary responsibility for the member's contributions. Though I maintain, for reasons provided in previous sections, that the radius of secondary responsibility is maximal, there is theoretical room for an account in which an individual does not bear secondary responsibility for some of the members.

I will focus on the *magnitude* of secondary responsibility, which is the degree to which to a given member bears responsibility for another given member's contributions. An individual might be fully responsible for the contributions made by someone else, or the individual might bear less than full responsibility. For example, one member's intentions might limit the magnitude of secondary responsibility for the contributions made by another particular member. Or one member might have intentions that limit the magnitude of secondary responsibility that other members bear. Since the magnitude of secondary responsibility any member bears is determined both by *her* intentions, as well as those of the other members of the collective, the magnitude of secondary responsibility a member bears is indexed to a given fellow member. So there is no single value for the magnitude of secondary responsibility that an individual bears (unless the individual is one of two members of a collective).

7.4.1. SHARED CONTRIBUTORY INTENTIONS

Shared intentions to contribute to a collective act result in significantly greater magnitude of secondary responsibility. In this section, I elucidate the notion of shared contributory intentions. I also distinguish between two types of shared contributory intentions, one of which results in greater secondary responsibility than the other. That is, I distinguish between cases in which the members who share contributory intentions also share a desire for the success of the object of the contributions, from those cases in which the members who share contributory intentions do *not* share a desire for the success of the

object of the contributions. The claims regarding the magnitude of secondary responsibility in this section are comparative and ordinal: shared contributory intentions with shared desires for the success of a collective act result in greater secondary responsibility than shared contributory intentions *without* shared desires for the success of the collective act, which, in turn, result in greater secondary responsibility than unshared contributory intentions.

Individuals have a shared contributory intention only if they intend to contribute to the object of their functional role in order to promote the achievement of that object. Moreover, for contributory intentions to be shared, the object of each contributory intention must be intended under the same description. So if I intend to contribute to the election of a Democrat, and you intend to contribute to the election of a Republican, then we do not share an intention to contribute to the election of the same person. Likewise, if I intend to contribute to the election of the tallest candidate and you intend to contribute to the election of the youngest candidate, then we do not share intentions to contribute to the election of the same person, even if the tallest person is also the youngest person. It is not enough that the object of the respective intended contributions is the same -- it must be intended under the same description as well.

Individuals in a cooperative collective might share contributory intentions only at a sufficient level of generalization. Take the previous example in which two individuals intend to contribute to the election of a Democrat and a Republican, respectively. Though they do not share a contributory intention with respect to whose candidacy to promote,

they do share an intention to contribute to the election of *some* candidate. This is a generalized contributory intention that they share. Thus, if an individual intends to contribute to the object of her functional role, then it is almost certain that she shared an intention to do so with other members, under *some* description, even if the description is “to contribute to a collective act”.

Sharing contributory intentions increases the magnitude of secondary responsibility. If members m_i and m_j share contributory intentions with respect to Σ , a shared object of their functional roles, then m_i and m_j bear greater secondary responsibility for one another's contributions to Σ than they would if they did not share contributory intentions with respect to Σ . As a result, members who contribute to the same *subsidiary* ends are in a position to bear greater secondary responsibility than those who contribute to the ultimate ends of a collective, since members who contribute to the same subsidiary ends are in a position to share more contributory intentions. If members m_i and m_j have functional roles that require contributions to ψ since ψ is instrumental to Σ , and if m_i and m_j share contributory intentions with respect to ψ and to Σ , then they bear greater secondary responsibility for one another's contributions not only to Σ but to ψ as well.

Often, sharing contributory intentions will involve desiring the success of a collective act. Consider a mechanic, m , specializing in the maintenance of ambulances. He performs his duties not merely to advance his own interests, but also to facilitate the collective act of saving the lives of all those who require urgent medical care from the hospital with which

he is associated. To the end, m intends to contribute to the overall maintenance of the fleet of ambulances operated by the hospital. Thus he bears secondary responsibility for the contributions made by others who have functional roles that require them to maintain the fleet of ambulances, and who also intend to contribute to the overall maintenance of the fleet of ambulances.

Consider, in addition, a nurse, n, working in an ER ward in the same hospital as the mechanic, m. Suppose n also performs his duties both to advance his own interests, and to facilitate the collective act of saving the lives of all those who require urgent medical care from the hospital where he works. Of course, the manner in which the *mechanic* participates in the collective act is quite different from the manner in which the nurse participates. But if they both have a coherent intention of contributing to the collective act of saving the lives of all those who require urgent medical care from the hospital, then they, under that sufficiently generalized description, share contributory intentions and thus bear greater secondary responsibility for one another's contributions when they contribute together to that shared object. The difference in the kind of help they provide is reflected in the absence of shared *subsidiary* intentions. Due to the difference in their jobs, the mechanic cannot coherently share with the nurse an intention to contribute to ends *subsidiary* to the general end of saving lives. However, the nurse can indeed share subsidiary intentions with her fellow nurses, such as the intention to administer life-saving medical intervention to patients who require it.

Shared contributory intentions need not be altruistic. Consider members of a small crime syndicate working together to rob a casino. Each member of the syndicate is motivated solely by the self-directed end of amassing wealth. Any individual contributing to the crime does so solely because she believes that her personal contribution will increase the chances of the crime's success. While each criminal's reasons for participating in the crime are agent-centered, each criminal still has a desire for the success of the crime -- not because the criminals care whether the others obtain their share of the wealth, but rather because, for any criminal, *ci*, the success of the crime is instrumental to *ci*'s share of the wealth. If *ci* is practically rational, she will want the collective end to be achieved. Because the achievement of the self-directed end requires the achievement of the collective act, *ci* has an instrumental reason to contribute to the collective end, since doing so is a means to achieving *ci*'s self-directed end. Such members share an intention to contribute to the success of the crime, and they desire the success of the crime.

Or, suppose, for example, that an Allied bomber in the Second World War volunteers to take part in the fire-bombing of Dresden, not just because he believes that his personal act will promote the success of the collective act, but also because he has the self-directed end to be among those who directly promote the success of the collective act. The pilot coherently intends to achieve an end with *de se* content -- to be among those who firebomb Dresden -- the achievement of which presupposes the success of the collective act. As a result, the success of the collective act is instrumental to the achievement of the pilot's self-directed end. If the pilot is practically rational, he will care whether the collective act occurs, because the success of the collective act is instrumental to the

achievement of his self-directed end. Thus the pilot is in a position to share contributory intentions with other pilots who similarly desire the success of the collective act, whether this desire is instrumental, ultimate, self-serving, or other-serving.

It is possible for a member of a collective to share contributory intentions with others, even if she does not desire the success of the object of the contributory intention. In such a case, the member bears less secondary responsibility for the contributions made by others, than she would if she desired or valued the achievement of the shared object. But why would an individual intend to contribute to a collective act the achievement of which she does not particularly desire?

Suppose that the motivating reason to contribute is not to promote the achievement of the collective act, but rather to promote the achievement of some other event caused by contributing to the collective act. In such a case, the member can coherently intend to contribute to a collective act, even if she does not desire the success of the collective act. For example, low-level employees of a corporation intentionally contribute to a collective act, not because they desire or value the collective act to which they contribute, but because contributing is a means to getting paid. That the employees are indifferent to the success of the collective act to which they intentionally contribute, does not preclude rationally intending to contribute to the collective act, since *the contribution* -- rather than the success of the collective act -- is instrumental to getting paid.

Members indifferent to the success of the collective act to which they contribute share an intention to contribute to that act if they intend to contribute to the collective act under the same description. In such a case, they share contributory intentions, even though none of them have any particular desire for the success of the collective acts to which they intentionally contribute. Any such individual bears greater secondary responsibility for the contributions made by others in the collective than she would if she did not share contributory intentions with them; but she bears less secondary responsibility for the contributions made by others in the collective than she would if she also shared a desire for the success of the collective act to which she intentionally contributes.

Here is another example in which a member shares with others a contributory intention, but does so without sharing a desire for the success of the collective act to which she intentionally contributes. Suppose a crime syndicate member participates in a crime not because she wants to amass wealth -- suppose she has enough already -- but because she wants to be admired by others for her willingness to participate in a dangerous venture. Since this is her sole motivation, she does not particularly care whether the crime is successful. She does not share with the others a desire for the success of the crime. However, she has an instrumental intention to contribute to the crime, and as a result shares this intention with others in the collective. The magnitude of secondary responsibility she bears is greater than it would be if she didn't share an intention to contribute to the crime. But it is not as great as it could be, since she does not share a *desire* for the success of crime.

Or consider again the example of the Allied bomber in the Second World War, but suppose he volunteers to take part in the fire-bombing of Dresden, not because he values or desires the success of the collective endeavor, but rather just for the thrill of killing helpless civilians. Again, the pilot coherently intends to contribute to the destruction of Dresden in order to achieve an end with *de se* content -- to experience the thrill of killing civilians -- but the achievement of this end does *not* presuppose the success of the collective act. Thus the pilot shares a contributory intention with other pilots who similarly intend to contribute to the destruction of Dresden; but he does not share with them any particular desire for the wholesale destruction of Dresden. As a result, he shares contributory intentions, but not a desire for the success of the collective act. Again, in such a case, where members do not share a desire for the achievement of the collective act to which they intentionally contribute, the members bear less secondary responsibility for one another's contributions to the shared object of their functional roles, than they would if they shared a desire for the achievement of the shared object.

In summary, any two contributing members, r_1 and r_2 , bear the greatest degree of secondary responsibility for the contributions made by one another if both members desire that Σ and share an intention to contribute to Σ . If both individuals share an intention to contribute to Σ , but only r_1 desires that Σ , then r_2 , who does not have a particular desire for Σ , will bear less secondary responsibility for the contribution made by r_1 . If r_1 and r_2 do not share an intention to contribute to Σ because r_2 does not have an intention to contribute to Σ , then r_2 will bear even less secondary responsibility for the contribution made by r_1 .

7.4.2. UNSHARED CONTRIBUTORY INTENTIONS

If a member has no particular intention to contribute to the object of her functional role, then, a fortiori, she has no shared intention to so contribute. As a result, she bears less secondary responsibility for the contributions made by others than would be the case if she did in fact share with others an intention to contribute. However, because she knows that she is contributing to the object of her functional role, she bears more secondary responsibility for the contributions made by others than would be the case if she did not know that that she is contributing.

Merely knowingly contributing to the object of a functional role diminishes secondary responsibility relative to intentionally contributing to the object of that functional role, for reasons similar to why, in the Model Penal Code, the fault element of an offence is stronger for *intention* than it is for *knowledge*. An agent who intentionally commits an offence is, ceteris paribus, more blameworthy than an agent who knowingly commits the same offense as a side-effect of some other intention. In the same way an individual bears greater blame for a wrong committed intentionally than she would if she committed the wrong merely knowingly, an individual bears greater secondary responsibility for the wrongs that *others* commit if she intentionally co-contributes to the wrong than she would if she merely knowingly did so. If I know that it is my purpose to contribute to a particular object of my functional role, and that it is the purpose of others to so contribute as well, and that we are together contributing, then I will bear, because of this knowledge,

some secondary responsibility for what others do in furtherance of the shared object of our respective functional roles. But an individual who intentionally contributes to a collective act is *more* responsible for the contributions made by others, than she would be if she non-intentionally contributed to that collective act.

An upshot is that the degree of secondary responsibility that a member bears is sensitive to her relevant mental states. A common-sense criticism of individual responsibility for collectively committed harms, is that it risks imposing on the members of a collective a blanket-judgment insensitive to intuitively relevant distinctions among these various individuals. That is, if we hold all the members of the collective responsible for a collectively committed harm, and in so doing we ignore key differences in intentions and beliefs among these individuals, then we fail to treat the members as autonomous individuals. On my account, though a basis for an individual's secondary responsibility is the choice to become the subject of a functional role, the individual's intentions and beliefs nonetheless affect her secondary responsibility. In what follows, I specify how various kinds of unintended contributions affect secondary responsibility.

7.4.3. RECKLESSNESS, NEGLIGENCE, AND ELLIPTICAL COMMITMENT

Recall that I am exploring factors that diminish secondary responsibility that a member bears for the contributions made by other members. I have already argued that secondary responsibility is diminished when the individual does not intend to contribute to the collective act under a description which other members intend to contribute. Secondary

responsibility for collectively committed harms is also diminished attenuated when the responsible individual contributes *recklessly* or *negligently*. I borrow the concepts of ‘recklessness’ and ‘negligence’ from criminal law and use them to elucidate the corresponding moral concepts. I then show how the moral concepts of recklessness and negligence affect secondary responsibility for collectively committed harms.

In criminal law, an act, ϕ , is negligent iff 1) it exposes others to a substantial risk of harm, 2) the individual who committed ϕ did not consider its potentially adverse consequences, 3) a reasonable person would have been aware of the risk resulting from committing ϕ , and 4) that reasonable person would have either refrained from committing ϕ or would taken measures that would have substantially reduced the risk resulting from committing ϕ . An act is *reckless* iff 1) it exposes others to a substantial risk of harm, 2) the individual who committed ϕ is aware of the potentially adverse consequences of committing ϕ , 3) the individual does not intend or desire the occurrence of these adverse consequences. The distinction between recklessness and negligence lies in the presence or absence of knowledge regarding the risk associated with either action. A negligent act in general incurs less culpability than a reckless act.

The moral concepts of recklessness and negligence, as they apply to individual contributions to collectively committed harms, are closely related to the legal concepts of recklessness and negligence. A member of a collective contributes to a collectively committed harm recklessly if she has no particular desire to contribute to the collective harm Σ , yet nonetheless decides to act according to her functional role, even though by so

doing she risks contributing to Σ . Such a member commits a *reckless contributory* act. If a member of a collective is *not* aware that she risks contributing to Σ by acting according to her functional role, even though she is in a position to know of the risk, then she commits a *negligent contributory* act. While a reckless contributor acts according to her functional role even though that she knows there is a significant risk that she is contributing to a harm, a negligent contributor is unaware of the significant risk that she is contributing to a harm. As with their analogues in private individual action, failure to recognize this risk is what distinguishes a negligent contributory act from a reckless contributory act. In cases of contributory recklessness, the secondary responsibility that the individual bears is only slightly less than the secondary responsibility that the individual would bear if she knowingly contributed to the impermissible collective act. And, of course, the secondary responsibility resulting from contributory recklessness is less than the secondary responsibility resulting from knowingly but unintentionally contributing to a collectively committed harm. And, in turn, the secondary responsibility resulting from knowingly but unintentionally contributing a collectively committed harm is less than the secondary responsibility resulting from intentionally contributing to a collectively committed harm.

It is important to distinguish between negligent or reckless *contributions*, from negligent or reckless *collective acts*. Here I am concerned solely with the former, and its effect on secondary responsibility. As an example of the latter, suppose m_i joins a street-racing gang; she contributes to illegal drag-racing that poses a substantial risk to those involved in the race as well as to others. Suppose m_i is aware of this. If an accident occurs and

others are unjustifiably hurt as a result of the drag-race, then mi contributed to a *reckless collective act*. As a result, she bears primary responsibility for her own contribution to this reckless collective act, and secondary responsibility for the contributions of others. It is certainly possible, however, for an individual's contribution to be reckless (or negligent), even if the object of her contribution is not a reckless (or negligent) act. For example, suppose a group plans on committing a murder; this collective act is not merely negligent or reckless -- it is intentional. Nonetheless, it is possible for a member, unaware of all the relevant non-moral facts, to contribute merely recklessly or negligently to this intentional collective act. It is this type of recklessness or negligence that I will investigate here.

Negligence, as I have claimed, involves ignorance of certain non-normative facts, viz., the risk associated with committing the negligent act. In the context of individual contributions to collective action, there are two kinds of non-normative ignorance that can diminish the degree of secondary responsibility that the contributors bear. The difference between the two involves the *object* of ignorance. A contributing member of a collective might not know what the object of her functional role is; that is, she might not know or might be uncertain regarding what collective act she is contributing to. Call this kind of ignorance 'object-based'.

Alternatively, a contributing member, though she knows what the object of her role is, might not know or might be uncertain regarding whether acting according to her functional role actually contributes to the collective act. Call this kind of ignorance,

‘contribution-based’. Either kind of ignorance affects whether the contributor is aware that by committing a particular act -- an act in accordance with her functional role -- she risks of contributing to a harmful collective act. We thus have two types of contributory negligence, correlating with contribution-based ignorance and object-based ignorance, respectively.

These types of ignorance -- object-based and contribution-based -- also describe types of *risk* associated with acting according to a functional role that contributes to a collectively committed harm. Suppose we know that the object of a functional role is an unjustified harm; but we cannot be certain that acting according to a functional role will result in a successful contribution to its object. If there is a substantial risk that acting according to the functional role will result in such a contribution, then the risk associated with acting according to the functional role *is contribution-based*. Suppose, alternatively, we know that acting according to a functional role will result in a successful contribution to its object; but we cannot be certain that the object constitutes a harm. If there is a substantial risk that the object is indeed an unjustified harm, then the risk associated with acting according to the functional role *is object-based*. Thus, along with the two types of contributory negligence, we have two types of contributory recklessness, correlating with contribution-based risk and object-based risk.

Since the distinction between contribution-based and object-based ignorance and risk cross-cuts the distinction between negligence and recklessness, we have four kinds of

contributions for which the contributor bears less secondary responsibility than she would bear if she contributed knowingly. I will discuss each in turn.

7.4.3.1. CONTRIBUTION-BASED NEGLIGENCE AND RECKLESSNESS

Suppose *ma* is an individual aware of the fact that a particular collective is committing harmful acts. She is also aware that she is a member of this collective. However, acting according to her functional role does not guarantee a successful contribution to the object of her functional role, i.e., the harmful collective act. Her functional role is so subsidiary, and so causally distant with respect to its ultimate object, that there is a good chance acting according to her functional role will not result in a contribution to the harmful collective act. Nonetheless, there is a substantial risk that if *ma* acts according to her functional role, this will result in a contribution to the harmful collective act. Moreover, a reasonable person would recognize that acting according to her functional role involves a substantial risk of a successful contribution -- after all, it is *ma*'s purpose to contribute to the object of her functional role, which she knows to be harmful. Though *ma* is aware that she is acting according to her functional role, and that her *purpose* in so doing is to contribute to a collective act, and that this collective act is harmful, she fails to recognize that by acting according to her functional role she runs a substantial risk of contributing to the collective act. As a result, if *ma* actually contributes to the object of her functional role, the contribution is negligent, where the object of her culpable ignorance is the likely causal link between acting as part of the collective, and contributing to the harmful collective act. This is an example of contribution-based negligence.

Contrast this with a case in which ma is in fact aware of the substantial risk that she is contributing to the harmful act by acting according to her functional role. Yet she nonetheless acts according to her functional role and as a result runs the risk of contributing to the harmful collective act, in spite of the fact that she has no intention to so contribute. Her contribution in such a case is reckless; the object of her culpable ignorance is, again, the likely causal link between acting as part of the collective, and contributing to the harmful collective act. This is an example of contribution-based recklessness.

7.4.3.2. OBJECT-BASED NEGLIGENCE AND RECKLESSNESS

In the previous section I discussed two kinds of contributions for which a contributor would bear less secondary responsibility than she would bear if she contributed knowingly, all things being equal. In what follows I discuss another two kinds of contributions for which a contributor would bear less secondary responsibility than she would bear if she contributed knowingly. The larger goal is to elucidate the factors that diminish secondary responsibility that a contributing member of a collective would bear for the contributions made by others.

Suppose mb knows that acting according to her functional role will result in a contribution to its object. The object of her functional role is likely to constitute an unjustified harm. However, mb is unaware of this risk -- she unreasonably and

mistakenly believes that the collective act to which she contributes is not an unjustified harm. She acts according to her functional role, and though she knowingly contributes to a collective act, she is unaware that this act is a harm. As a result, she unknowingly risks contributing to a collectively committed harm, in spite of the fact that she has no particular desire to contribute to such a harm. This is an example of object-based negligence. Because her ignorance of this risk is culpable, she bears secondary responsibility for the contributions of others, albeit less than she would bear if her contribution were reckless instead of negligent.

Alternatively, suppose mb knows that acting according to her functional role will result in a contribution to its object, and she knows that there is a substantial risk that the object of her functional role is a collectively committed harm. She does not intend to contribute to this risk. She acts according to her functional role nonetheless, ignoring the risk of contributing to the harmful collective act. As a result, her contribution is reckless, where the basis of the recklessness lies in the risk that the object of her contribution is an unjustified harm. This is an example of object-based recklessness. She bears greater secondary responsibility for the contributions of others than she would if her contribution were merely negligent.

7.4.3.3. COMBINING OBJECT-BASED AND CONTRIBUTION-BASED RISK

Object-based recklessness does not, in general, result in less or more secondary responsibility than contribution-based recklessness. And object-based negligence does

not, in general, result in less or more secondary responsibility than contribution-based recklessness. However, object-based and contribution-based recklessness or negligence can be combined, resulting in a contribution that is 1) doubly reckless, 2) doubly negligent, or 3) reckless and negligent.

In (1), contribution-based recklessness is combined with object-based recklessness. The member acts according to her functional role, even though she is aware that by so doing she risks contributing to a collective act, and even though she is aware of the substantial likelihood that the collective act is harmful. The resulting secondary responsibility in such a case is, in general, less than the secondary responsibility resulting from a contribution that is both reckless and negligent, i.e., (2). If a contribution is both reckless and negligent, then the member is culpably unaware either that by acting according to her functional role she risks contributing to a collective act, or that the collective act is likely to be harmful. If the contribution is both reckless and negligent, and if she is culpably unaware that by acting according to her functional role she risks contributing to a collective act, then she is aware that the collective act is likely to be harmful. This is an example of contribution-based negligence and object-based recklessness. If the contribution is both reckless and negligent, and if she is culpably unaware that the collective act is likely to be harmful, then she is aware that by acting according to her functional role she risks contributing to a collective act. This is an example of contribution-based recklessness and object-based negligence. Though either combination of recklessness and negligence incurs, in general, less secondary responsibility than object/contribution-based recklessness/negligence alone, a contribution that is both

reckless and negligent still incurs greater secondary responsibility than that which results from a doubly negligent contribution, i.e., (3). In a doubly negligent contribution, the member is culpably unaware that by acting according to her functional role she risks contributing to a collective act, and she is culpably unaware that the collective act is likely to be harmful.

7.4.3.4. RECKLESS CONTRIBUTIONS AND ELLIPTICAL COMMITMENT

If an individual agrees to commit certain tasks without an intention to contribute to the collective act, then the individual has *elliptically* committed herself to the contributing to the collective. In this section I distinguish reckless contributions from contributions made by a member who has elliptically committed herself to contributing to the object of her functional role.

Recall that an individual can be a member of a collective, even if she agrees only to follow a particular set of rules. Since intentions are not closed under logical (let alone causal) implication, the fact that she intentionally agrees to act according to a particular set of rules does not entail that she intends to contribute to the object of her functional role. I argued that this individual has nonetheless committed herself to so contributing as a result of agreeing to abide by a set of rules which she is in a position to recognize as contributory to a collective act. More completely, she has committed herself to providing and accepting assistance with respect to the accomplishment of the object of her

functional role. As a result, the individual bears secondary responsibility for the contributions of others in the collective who act qua members of the collective.

Suppose the object of her functional role is a wrongful act. If the member has no intention to contribute to the collective act, and has not explicitly agreed to assist or to accept assistance in furtherance of the collective act (i.e., the object of her functional role), but has agreed to perform a task which she recognizes as instrumental to a wrongful act, then she might seem to be contributing recklessly. But this is not so. For a contribution to be reckless, the contributor must be less than certain that her contribution is a contribution to a collectively committed harm. The uncertainty might be object-based, in which case the harmfulness of the collectively committed act is in question, or the uncertainty might be contribution-based, in which case the efficacy of the contribution is in question. As in reckless contributions, the individual agrees to commit certain tasks without an intention to contribute to the collective act. But unlike those who contribute recklessly, a member whose contribution results from an elliptical commitment can be certain that by acting according to her functional role she is contributing to a collectively committed harm.

7.5. DEGREE OF PRIMARY RESPONSIBILITY

In the previous sections, I discussed how various factors affect the degree of secondary responsibility that an individual member bears for the contributions made by other members. I will discuss how various factors affect the degree of primary responsibility.

But a comprehensive account of what I have been calling ‘primary responsibility’ has been a goal in Normative Ethics for a very long time. An account of primary responsibility is, after all, an account of individual responsibility for one’s own actions. The account that I am developing -- an account of individual responsibility for collectively committed harms -- is meant to be compatible with competing theories of primary responsibility. As a result, my claims regarding what determines the degree of primary responsibility that an individual bears will be limited to those that I take to be necessary given the account of secondary responsibility that I have developed so far.

Broadly generalized, I focus on two factors that determine an individual’s primary responsibility. The first factor is the *object* of responsibility. In the next section, I explore three objects for which a member might bear primary responsibility. The second factor is the *method* by which the member contributes to the collective act. This determines not what the object of responsibility is, but rather the degree to which an individual bears primary responsibility for the object of her contribution.

7.5.1. THE OBJECT OF PRIMARY RESPONSIBILITY

Recall that what separates primary from secondary responsibility is the object of responsibility. An agent bears primary responsibility for her own actions, rather than for the actions of others. For example, suppose five individuals are together building a bomb. Each individual brings fuel for the bomb -- some bring more fuel than others. The more fuel is added to the bomb, the bigger its explosion will be. (To keep things simple,

suppose the relation between the amount of fuel and explosive power is linear, as is the relation between the explosive power and the amount of destruction caused by the bomb). Each member will bear primary responsibility for the increase in destruction caused by the fuel she adds to the bomb. No member will bear primary responsibility for more than what she causes. So if m1 adds more fuel than another member, m2, then m1 causes a larger portion of the destruction wrought by the bomb. Accordingly, the primary responsibility that m1 bears is greater than the primary responsibility that m2 bears; this is because the object of primary responsibility is the portion of the destruction that she causes.

This is little more than an act-consequentialist analysis of primary responsibility. Does this mean, then, that primary responsibility disappears, in the manner described in section 6.1.1? After all, part of the purpose of secondary responsibility was to explain how responsibility remains intact when a member provides an overdetermined contribution to a collective act. In the next few sections, I will argue that primary responsibility for a contribution does not completely disappear, even when the contribution is overdetermined. I will argue that certain kinds of overdetermined contributions affect the *propensity* of a collective act's occurrence. In such cases, the object of the contributor's primary responsibility is not the collective act, but the propensity of the collective act's occurrence. I will explain what I mean by 'propensity' and how it differs from 'probability'. But before I explain how the propensity of a collective act can be the object of an overdetermining contribution to that collective act, I will distinguish two types of overdetermining contributions.

7.5.1.1. OVERDETERMINATION AND PRIMARY RESPONSIBILITY

All things being equal, a member of a collective bears primary responsibility for the collective act itself if the collective act would not have occurred had the member refrained from contributing to the collective act. In such a case, the collective act itself is the object of responsibility. It might be tempting to claim that the collective act is the object of responsibility because the contributor *caused* the collective act. But such a claim is misleading, since a member can cause a collective act by contributing to it even if the collective act would have occurred had the member refrained from contributing. It is useful, then, to distinguish two types of causation in order to be clear regarding the conditions under which a member who contributes to a collective act is responsible for the collective act as a whole. One type of causation requires the occurrence of its object, while the other type does not. Suppose that given the absence of the contribution ψ , Σ would not occur. In such a case, ψ can be called a “but-for” cause of Σ since there is a dependence-relation between ψ and Σ . Suppose, alternatively, that a member’s contribution ψ *would* cause Σ if it were not for another member who contributes to Σ by committing ϕ , which preempts ψ by causing Σ first. As a result, the occurrence of Σ does not depend on the occurrence of ϕ . In such a case, ϕ can be called a “preemptively overdetermining” cause of Σ , since Σ is overdetermined as a result of the actual causal influence that ϕ has on Σ , and the counterfactual causal influence that ψ has on Σ .

As an example, suppose a killer, k1, shoots d in the head, resulting in d's death. Another killer, k2, shoots d in the head one second after k1 does so. The death of d is overdetermined, in that if k1 had not killed d, then k2 would have done so. But as it is, k2's bullet strikes d after d has been killed by k1. So k1 causes d's death, in that the bullet's physical interaction with d (in combination with background conditions, such as the laws of physics) is sufficient for d's death. But k1's act preemptively overdetermines the death of d, since d would have been killed one second later anyway. As a result, k1's act is a preemptively overdetermining cause of d's death. The murder k1 commits preempts the murder that k2 would have committed in the place of k1, had k1 missed the target. So the death of the target is overdetermined by the acts committed by k1 and k2. Suppose, however, that k1 misses; k2, then, successfully kills d. In this case, since d would not have been killed were it not for k2, k2 is a *but-for* cause of d's death, rather than a preemptively overdetermining cause of d's death.

It might be argued that the apparent distinction between but-for causes and preemptively overdetermining causes is specious, resulting from a coarse-grained reference to the object of causation, which in the previous example is d's death. If we drop this reference in favor of the more fine-grained distinction between death-by-k1 and death-by-k2, then there is no relevant event overdetermined by the actions of k1 and k2. As a result, k1's act is a but-for cause of death-by-k1; were it not for k1's act, then death-by-k1 would not have occurred. Similarly, k2's act is a but-for cause of death-by-k2; were it not for k2's act, then death-by-k2 would not have occurred. On this analysis, there is no preemptively

overdetermining causation. I won't belabor criticizing this view, as it is has been adequately done so by David Lewis.

A contribution that consists of a preemptively overdetermining cause can be called a 'preemptively overdetermining contribution'. And a contribution that consists of a but-for cause, can be called a "but-for contribution". If a contribution is a but-for contribution, then the object of responsibility is the collective act as a whole, and the member who provides the strong contribution will bear primary responsibility for that collective act as a whole. This is considerable primary responsibility, but it is appropriate given that without the member's (but-for) contribution, the collective act would not have occurred at all.

Contrast the example of the overdetermined murder with a nearby case which demonstrates a different kind of overdetermination. Suppose k_1 , k_2 , and k_3 together beat d to death. If any one of k_1 , k_2 , or k_3 had not participated, d would have died as a result of the assault inflicted by the other two. However, if any *two* had not participated, d would have lived. The death of d is overdetermined by what k_1 , k_2 , and k_3 do, but not because some of the members are disposed to act in place of those who fail to kill d . Rather, the death of d is overdetermined by what k_1 , k_2 , and k_3 do, because any two of the contributions (in combination with the relevant background conditions) are together sufficient for the death of d , even though all three of the contributions severally contribute to the death of d . This type of overdetermination can be called "superfluous overdetermination"; and a contribution partly in virtue of which a collective act is

superfluously overdetermined can be called a “superfluously overdetermining contribution”. I discuss primary responsibility associated with preemptively and superfluously overdetermining contributions in the next two sections.

7.5.1.2. SUPERFLUOUS OVERDETERMINATION & PRIMARY RESPONSIBILITY

A superfluously overdetermining contribution to Σ , like a preemptively overdetermining contribution to Σ , is unnecessary for the occurrence of Σ in the sense that in the absence of the contribution, Σ would nonetheless occur. But unlike a preemptively overdetermining contribution, a superfluously overdetermining contribution does not cause Σ -- rather, it is part of what causes Σ .

Suppose a dozen individuals, d1 through d12 together decide to push a boulder off a cliff. For any individual, the effort put forth by all the others is enough to push the boulder off the cliff. So for any individual, the contribution she provides is superfluously overdetermining, not merely because the boulder would be pushed off the cliff in her absence, but also because she actually contributes to pushing the boulder off the cliff. Contrast this with the following case. Suppose that d6 generates so much force against the boulder that his contribution alone is sufficient to push the boulder off the cliff, and that without his contribution the others would fail to push the boulder over the cliff. However, unlike the others, d6 has a back-up -- an individual as strong as d6 is ready to take d6's place, should the latter fail to contribute. Thus d6 provides a preemptively overdetermining contribution to the collective act.

I have claimed that the object of a member's primary responsibility for providing a but-for contribution to a collective act is the collective act itself. When a member provides a superfluously overdetermining contribution to a collective act, the object of responsibility is *not* the collective act to which she contributes. Nor is her contribution to the collective act an object of her primary responsibility, because, again, the effects of her contribution are overdetermined. However, superfluously overdetermining contributions to a collective act can result in primary responsibility for increasing the *propensity* of the collective act's occurrence. This increase in the propensity of the collective act's occurrence is not overdetermined, even if it is the result of a superfluously overdetermining contribution. But what does it mean to increase the propensity of a collective act's occurrence?

The propensity of an event, like the frequency of an event, and unlike the subjective probability of an event's occurrence, is not a mental or logical abstraction, but is instead a physical property. More specifically, it is a disposition of a particular initial event or state of affairs to yield a particular outcome. This disposition is not a subjective assessment of probability, but is instead the physical basis for the set of outcomes that would occur if the initial event were 'repeated' many times over. For example, suppose that a particular collective act, Σ , requires at least 30 contributors. The propensity of 30 contributors to achieve Σ is high, but not as high as the propensity of, for example, 40 contributors. In the former case, if but one contributor fails to act according to his functional role, then Σ will fail. By including an additional ten contributors, we have created a 'safety-net',

which allows up to 10 of the contributors to fail, or to contribute less effectively than they normally would, without sacrificing Σ . If the initial conditions consisting of 40 contributors were instantiated many times, Σ would successfully occur more often than would be the case if the initial conditions consisting of 30 contributors were instantiated many times. This is just to say that 40 contributors have a higher propensity of achieving Σ . (Obviously, diminishing returns with respect to the increase in propensity would set in as the number of contributors increases).

Consider again the previous example; d12 provides a superfluously overdetermining contribution, ϕ , to the collective act of Σ , i.e. pushing the boulder off the cliff. Even though his contribution is superfluously overdetermining, it nonetheless increases the propensity of successfully pushing the boulder off the cliff, since his contribution would make the difference in cases where some other contributor fails to perform as well as he or she normally would. Likewise, suppose another contributor, d11, commits ψ , which is a superfluously overdetermining contribution to pushing the boulder off the cliff. Like ϕ , the occurrence of ψ positively affects the propensity of successfully pushing the boulder off the cliff. But suppose that ψ positively affects the propensity of Σ 's occurrence to a greater degree than ϕ does -- perhaps d11 is stronger than d12. The object of d11's responsibility is his increase in the propensity of Σ 's occurrence. And the object of d12's responsibility is his increase in the propensity of Σ 's occurrence. As a result, even though d12 and d11 (severally) commit superfluously overdetermining contributions to Σ , d11 bears responsibility for more than d12 does, since the former increases the propensity of Σ 's occurrence to a greater degree than the latter does. The

individuals do not bear primary responsibility for Σ ; rather, each individual bears responsibility for increasing the propensity of Σ .

7.5.1.3. PREEMPTIVE OVERDETERMINATION AND PRIMARY RESPONSIBILITY

A member who provides a preemptively overdetermining contribution to Σ actually causes Σ -- the contribution is preemptively overdetermining in that Σ would have occurred even if the member had not caused Σ . But the member who provides a preemptively overdetermining contribution to Σ does not have Σ as an object of primary responsibility -- nor does she have her contribution as an object of primary responsibility. As in cases of superfluously overdetermining contributions, the object of primary responsibility is the propensity of Σ 's occurrence. However, preemptively overdetermining contributions affect the propensity of Σ in a manner different from the way that superfluously overdetermining contributions affect the propensity of Σ . Here I briefly investigate how preemptively overdetermining contributions to a collective act uniquely affect the propensity of a collective act.

Recall from section 4.1 that functional roles have eligibility-conditions which a candidate member must satisfy in order to be an actual member. The eligibility-conditions of a functional role affect the propensity of the occurrence of a role's object, even when the collective act is overdetermined, and especially when the collective act is preemptively overdetermined.

Functional roles vary with respect to the specificity of their eligibility requirements for a candidate subject-particular. For example, suppose a particular role requires skills so advanced, that only a handful of persons in the world have the requisite skills. Acting according to a functional role with eligibility criteria that are satisfiable by few persons will increase the propensity of the collective act to which the role is related, relative to a functional role with more general eligibility criteria. This is because the probability of finding subject-particulars for the latter is, *ceteris paribus*, greater than the probability of finding subject-particulars for the former.

If an individual eligible to become the subject-particular of a role with specific eligibility criteria declines to do so, then, *ceteris paribus*, the probability of filling the role diminishes to a greater degree than it would if an individual eligible to become the subject-particular of a *less* specific role declines to do so. Since a functional role is designed to be productive with respect to its object, if a functional role goes unfulfilled, then its object does not have a contribution that it would otherwise have. This decreases the propensity of the collective act's occurrence, even if the role provides only preemptively overdetermining contributions. This is due to the same reasons why superfluously overdetermining contributions can increase the propensity of a collective act's occurrence.

For example, suppose *j* is a member who commits ϕ , which is a preemptively overdetermining contribution prescribed by *j*'s functional role. If *j* had not been a member and had thus not committed ϕ , then someone else would have taken his place

and consequently would have committed ϕ . But suppose that though it is very likely that someone else would have taken j 's place, it is not guaranteed. So if j had not been a member, the propensity of Σ 's occurrence would have been lower, which is to say that the initial conditions more frequently yield Σ when j is a member, than when j is not a member. This is to say that j 's membership and consequent contribution to Σ increases the propensity of Σ . Suppose k is a member of the same collective as j ; and suppose that k makes the same contribution to Σ as j does. However, k 's functional role has eligibility criteria that are even more difficult to satisfy than j 's functional role. That is, if k had not been a member, it would have been harder to find a replacement for her, than it would have been to find a replacement for j had *she* not been a member. So though both j and k make the same preemptively overdetermining contributions to Σ , j 's contribution is more significant, since it increases the propensity of Σ 's occurrence to a greater degree than j 's membership does.

Since eligibility-conditions affect the propensity of a collective act's occurrence, the eligibility-conditions also affect the object of primary responsibility. A member who provides a contribution prescribed by a functional role with highly specific eligibility criteria will, *ceteris paribus*, contribute more to the propensity of the collective act's occurrence, and thus bear primary responsibility for a larger increase in the propensity of that act's occurrence, than will a member who provides a contribution prescribed by a functional role with less specific eligibility criteria.

I have argued that a member, j , who provides a preemptively overdetermining contribution to Σ increases the propensity of Σ -- but the member who acts as j 's 'back-up' also increases the propensity of Σ . This member, j 's back-up, is part of that in virtue of which Σ is preemptively overdetermined -- but the contribution that j 's back-up provides is wholly dispositional rather than actual. A dispositional contribution is productive only given the appropriate manifestation conditions. If the manifestation conditions do not occur, then the subject of the dispositional functional role provides no contribution to the collective act. But the subject-particular is nonetheless in a position to contribute to the collective act, should the manifestation conditions occur. Contrast this case with one in which the functional role is left 'unmanned' -- it has no subject-particular. In such a case, if the manifestation conditions of the functional role's contribution occur, there is no one to contribute to the collective act. So if the functional role, prescribing a dispositional contribution, simply has a subject-particular, then the propensity of the collective act's occurrence increases. That is, often, simply *filling* a functional role increases the propensity of a collective act.

As an example, suppose an understudy for the starring role of a theatrical production has a functional role that is only dispositionally productive. While the actual starring role in the play provides an actual contribution to the collective act of performing a play, the understudy contributes only if the starring role fails to contribute. But even if these manifestation conditions do not occur, the understudy, as the subject-particular of a functional role that contributes dispositionally, increases the propensity of the collective act's occurrence merely in virtue of being disposed to contribute. As a result, the

understudy bears primary responsibility for increasing the propensity of the collective act's occurrence, even though the understudy's contribution to the collective act is wholly dispositional.

But suppose that the play is a flop because the lead actor's performance was terrible -- and suppose the understudy would have done much better. Is it the case that understudy, by taking on her functional role, has increasing the propensity that the play is a flop? And if so, does the understudy bear primary responsibility for increasing the propensity that the play is a flop? This seems counterintuitive since the understudy's dispositional contribution is not what makes the play a flop. In this case, the understudy's participation does indeed increase the propensity that the play is a flop, since the understudy's participation increases the propensity that the play will be performed with the lead actor. But the understudy also increases the propensity that the play will be a success, since the understudy is a good performer; in those instances in which the understudy's dispositional contribution becomes actual, the play is a success. Both propensities have to be taken into consideration when we want to ascertain for what the understudy bears primary responsibility.

I have argued that individuals who provide preemptively overdetermining contributions have as an object of responsibility an increase in the propensity of a collective act. However, I have said nothing about *how much* responsibility is incurred in such cases. How culpable is an individual who increases the propensity of a morally bad collective act? And to what is that individual morally liable? I do not here provide specific answers

to these questions. Presumably, responsibility for a very high increase in the propensity of Σ is tantamount to responsibility for Σ itself. A small increase in the propensity of Σ is tantamount to little responsibility for Σ itself.

It should be noted, though, that according to modern criminal law, an individual who provides a preemptively overdetermining contribution to a collective act in the absence of mitigating circumstances bears full responsibility (or culpability, to be precise) for the collective act itself, rather than for increasing the propensity of its occurrence. Consider a modification of the hit-man example used in the previous section. Suppose that $k1^*$ shoots d ; if she had missed, d would have lived. In this case $k1^*$ bears primary responsibility for d 's death. In the original example, $k1$ and $k2$ work together in order to kill d ; if $k1$ misses, $k2$'s bullet would have killed d . Thus $k1$'s contribution is preemptively overdetermining. In criminal law, $k1$ and $k1^*$ would bear the same culpability, all other things being equal. On my account, $k1^*$ bears primary responsibility for d 's death, while $k1$ does not. At best, $k1$ bears primary responsibility for increasing the propensity of d 's death. In any case, on my account, $k1^*$ does *not* bear greater responsibility for d 's death than $k1$ does. After all, to determine $k1^*$'s responsibility, we need to include secondary responsibility as well. Once this is done, the results of my account accord with criminal law in the case described -- i.e., $k1^*$ and $k1$ are equally responsible, and for the same thing -- viz. for d 's death. On my account, $k1$ does *not* bear responsibility for d 's death, only if we ignore secondary responsibility -- which is what I have been doing here in order to explore what determines the degree of primary responsibility that contributors bear.

7.5.2. THE SUBSIDIARITY OF CONTRIBUTIONS

To reiterate, I am focusing here on factors affecting the degree of primary responsibility that a member bears for her own contributions to a collectively committed act. Settling the object of primary responsibility does not alone settle the degree of primary responsibility that the contributor bears. For example, whether a contribution consists of an *action* or an *omission* might affect how much primary responsibility the contributor bears. Here, I will explore how the *subsidiarity* of a contribution affects the degree to which an individual bears primary responsibility for the object of her contribution.

The subsidiarity of a contribution depends on how *mediated* that contribution is with respect to the collective act. But what is meant by ‘subsidiary’ or ‘mediated’ contributions? Suppose for a collective act Σ , there are intermediate events i_1 through i_n which are independently or jointly instrumentally related to Σ . Suppose that for one of these events, i_1 , there are, in turn, intermediate events i_{11} through i_{1n} which are independently or jointly instrumentally related to i_1 . Call these second-order instrumental events, ‘subsidiary’, where the term is transitive. So if there are events subsidiary to i_{11} , they will be subsidiary to i_1 , and to Σ . One event can be said to be ‘more’ subsidiary than another if the former is imbedded in a greater number of nested, instrumentally-related events than the latter. However, if one event is subsidiary to another, this does not mean that the former is necessary for the latter. Instead, it only means that the former is instrumentally related to the latter.

The degree of subsidiarity sometimes affects primary responsibility, in that a contribution with a higher degree of subsidiary might result in less primary responsibility than a lower degree of subsidiary. For example, in a hierarchical command structure, the most subordinate functional roles will, *ceteris paribus*, prescribe contributions with greater subsidiarity. This is because the contributions governed by subordinate functional roles contribute to Σ only via the contributions governed by functional roles that are less subordinate. Thus the contributions governed by subordinate functional roles are more subsidiary to Σ , and thus bear a greater degree of subsidiarity than the contributions governed by less subordinate functional roles. Obvious examples include relations of authority in a military context; but subsidiarity is integral to the functional roles in many collectives, beyond military ones. For instance, secretaries tend to have much more subsidiary functional roles in a corporation than the members of the Board of Directors have.

Though highly subsidiary contributions tend to be overdetermining with respect to the contribution's object, this is not always the case. It is possible for the contributions made by the subject-particulars of highly subsidiary functional roles to be but-for causes of the collective act. Suppose the boss of a crime syndicate decides to use a low-level newly-recruited thug in the robbery of a high-profile casino. The new recruit is the police chief's son; the mob boss knows that the chief-of-police will not authorize deadly force against his own son. Though the new recruit occupies the lowest rung in the chain of command among those robbing the casino, the contribution he provides is a but-for cause of the

robbery, in that the robbery would not be attempted without his participation. The functional role, then, is quite important since the contribution it prescribes is a but-for cause -- but it is likely that it is not as important as it would be if his contribution were less subsidiary.

Why believe that the subsidiarity of a contribution affects primary responsibility?

Primary responsibility tends to be conserved efficiently along causally related events. If j knows that causing event α will cause event β , which causes event γ ... and so on through ω , and if j intentionally commits α in order to cause ω , then typically j is as responsible for α as she is for ω . But consider a nearby case, in which j^* causes event α^* , which causes event β^* , and so on through ω^* . In this case, events β through ω are acts committed by distinct and fully autonomous individuals. Thus j commits α , knowing that this would cause someone else to commit β^* , which will cause someone else to commit γ^* , and so on through ω^* . Assume that j is *as certain* (and is as justified in her certainty) that committing α will ultimately cause ω , as j^* is certain that committing α^* will ultimately cause ω^* . Nonetheless, because in the latter case the relevant causal chain is mediated by autonomous agents, the responsibility that j^* bears for ω^* is, *ceteris paribus*, less than the responsibility that j bears for ω . We often do not regard primary responsibility as well-conserved along causally related events when some of those events are acts committed by other fully autonomous individuals. This intuition drives the claim that contributions from subsidiary functional roles often result in less primary responsibility than contributions from functional roles that are less subsidiary, even if

both contributions have the same effects on the collective act (or the propensity of the collective act's occurrence).

However, I do not mean to claim that a causal chain mediated by autonomous agent *always* diminishes the primary responsibility of the individual who set off the chain. Perhaps there are only certain kinds of causal relations between the individual setting off the chain and the mediating agent that diminishes the former's responsibility for what she causes. In any case, my intention here is not to present a comprehensive account of how subsidiarity affects primary responsibility. Rather, I only claim that often subsidiarity can and does affect primary responsibility -- and that an account of primary responsibility should be sensitive to this fact.

Consider a modification of Derek Parfit's example of the "Harmless Torturers" (1986). Suppose individuals are paid to participate in the torture of an innocent. Each individual, by participating in the torture, agrees to administer a small electric shock to the victim by pressing a button. In return, the administering participant is paid a fee. The electric shock that each individual administers is so small that it has no perceptible effect on the victim. However, the cumulative affect of the imperceptible shocks is extremely painful for the victim. Thus, while each individual commits an act that has an imperceptible effect on the victim, the collective itself commits torture. The individual contributions are overdetermined, in the sense that for any individual, if she opts not to contribute, there are more than enough other individuals whose participation will ensure not only that that the victim is tortured, but that the experience of it will be the same regardless of whether

the individual contributes to the torture. Clearly, then, an individual member's contribution to the collective act of torturing the victim is superfluously overdetermining, and only slightly affects the propensity of the collective act. Nonetheless, the individual's contribution is not subsidiary with respect to the contributions of others; each individual contributes directly to the harmful collective act. As a result, each individual's contribution has a very low degree of subsidiarity, which results in greater primary responsibility for the increase in the propensity of the collective act's occurrence. (Of course, members will, in addition, bear secondary responsibility for one another's contributions).

It should not be assumed that a highly subsidiary contribution necessarily constitutes only a small part of a collective act. For example, suppose a mixed aircraft and ship formation on active duty during a war is charged with the destruction of an enemy airbase. To accomplish this, enemy surface-to-air missile sites and fire control radars need to be destroyed first. Surface-to-surface cruise missiles are used to destroy these targets, before several aerial squadrons begin bombardment of the remaining hardened targets. Suppose that though destroying the enemy surface-to-air missile sites and fire control radars using cruise missiles is a subsidiary contribution to the goal of destroying the airbase, the destruction of the sites constitutes a substantial contribution, in its own right, to the destruction of the airbase. As a result, though the contributions consisting in the operation of the cruise missiles are more subsidiary relative to the contributions consisting in piloting the aircraft conducting the aerial bombardment, the former's contribution to the destruction of the airbase can be on par with if not greater than the latter's contribution.

I have explored three kinds of objects that contributions can have -- a) the effect a contribution has on a collective act, b) the collective act *tout court*, and c) the propensity of a collective act. I have also explored how the subsidiarity of a contribution can affect the degree of responsibility that an individual bears for a, b, or c. But this investigation of the factors determining the degree of primary responsibility that an individual contributor bears, is not meant to be exhaustive. Clearly, there are morally relevant important aspects of contributions that I am not covering here.

7.6. MITIGATED RESPONSIBILITY

In sections 7.4 and 7.5 I discussed how various factors affect -- or more specifically, diminish -- a member's responsibility in a collective. But none of the factors I have discussed involve *excuses*. There are a variety of excuses that mitigate a member's primary or secondary responsibility. I will not belabor explanations of all the specific types of mitigating circumstances; rather, I will review a few by showing how they apply in collective contexts.

7.6.1. NON-CULPABLE IGNORANCE

The most common type of excuse arises from non-culpable ignorance. An individual who implicitly agrees to be responsible for what the collective does might be reasonably unaware of or misinformed about any number of facts that affect primary or secondary

responsibility. The degree of responsibility the member bears in these cases depends on our views regarding the extent to which non-culpable ignorance mitigates responsibility. Ignorance regarding the pertinent moral and non-moral facts is unlikely to serve as an excuse if the ignorance is the agent's fault, in the relevant sense (as discussed in section 7.4.3). But if no rational person could have known the pertinent facts, then the member's responsibility (whether primary or secondary) is significantly mitigated. That is, her non-culpable ignorance serves as an excuse, diminishing or eliminating her responsibility for both her own contributions to the object of her functional role, as well as for the contributions made by other members.

Object-based and contribution-based ignorance (see section 7.4.3) might be culpable, as is the case when a contribution is reckless or negligent. Non-culpable object-based and contribution-based ignorance can serve as potential excuses mitigating primary and secondary responsibility. But in addition, there is a third type of ignorance that can serve as an excuse. Members of a collective bear secondary responsibility for the contributions made by others; but if a member, *m*, is non-culpably unaware that others in the collective, by acting according to their functional roles, are contributing harmfully, then this ignorance might insulate *m* from secondary responsibility for the harmful contributions made by the other members. This type of ignorance can be called 'member-based' ignorance. As with other forms of ignorance, it can be culpable or non-culpable.

There are a variety of ways that an individual might come to be non-culpably ignorant of the pertinent moral or non-moral facts in virtue of which her contribution to the collective

is unjustified. Consider the following type of non-culpable member-based ignorance resulting from deceit. Suppose *j* is a member of a collective; *j*'s functional role is morally benign. But included in the collective are members with functional roles requiring the subject-particular to commit harmful acts. These functional roles, however, are well hidden; the subject-particulars of these functional roles as well as a few influential members of the collective deceive the other members by both failing to disclose the existence of these functional roles and by concealing any indications of their existence. Moreover, there is no reason to believe that the collective is of a type disposed to include these harmful functional roles. Since any member bears secondary responsibility for the contributions made by all the other members, any member would bear secondary responsibility for the contributions made by the subject-particulars of the harmful functional roles. But the members are not to blame for their ignorance of the fact that the collective includes those harmful functional roles. This serves as an excuse, mitigating the secondary responsibility they bear for the contributions made by the subject-particulars of the harmful functional roles. Note that this excuse does not mitigate *primary* responsibility, since members do not bear primary responsibility for the act committed by other members.

When a member is non-culpably ignorant as a result of deceit, the deceivers need not be members of the collective. It could be other collectives or society at large that deceives the members in ways that mitigate the latter's responsibility for contributions to the collectively committed harm.

A member might be non-culpably ignorant as a result not of deception but of unfortunate happenstance. In such a case, there are morally relevant non-moral facts of which the member is ignorant, but non-culpably so, in that a reasonable person would be similarly ignorant given the circumstances. There is no deceit or manipulation in such cases; instead, there are morally relevant unforeseen events or states of affairs the occurrence of which were not reasonably foreseeable. For instance, suppose there is good reason to believe that the commission of a particular collective act, Σ , is not only morally permissible, but beneficial. If particular members have information that others do not have -- information which reveals that a risk of significant harm caused by Σ is higher than what has been reasonably foreseen -- then the primary and secondary responsibility of the members privy to this extra information are not mitigated (or at least, not to the extent that the responsibility of the other members is mitigated) since they are not ignorant of the relevant facts. The responsibility of those members who are unaware that a risk of significant harm caused by Σ is higher than what has been reasonably foreseen would indeed be mitigated.

7.6.2. COERCION

Coercion is another type of excuse that often motivates participation in collectives that commit harmful acts. In the same way that non-culpable ignorance can be either circumstantially caused or intentionally caused by a deceitful party, coerced participation in a collective can be either circumstantially caused or intentionally caused by a coercive party. Participation in a collective is intentionally coerced if the participation is the result

of a threat made by another party. The threatening party might be part of the collective of which the member is coerced into joining, or the threat might be from a third party. In either case, physical coercive threats can qualify as an excuse mitigating primary and secondary responsibility in a collective if 1) the threat is of serious bodily harm or death, 2) the threat is otherwise unavoidable, and 3) the threatened party is not responsible for being threatened. Physically coercive threats used to motivate participation in a collective are sometimes legal; for example, during a military draft the threat of incarceration leveled against candidate conscripts is used to coerce them into joining the military. This kind of coercive threat is often cited as an excuse for participation in an unjust war.

In addition to threats of physical harm, there are threats that operate by imposing psychological stress, and threats of nonphysical harms, such as harms to a person's reputation or wealth. Typically, the latter kind of threat, of which blackmail and extortion are examples, is less likely to completely mitigate primary or secondary responsibility, since such threats are in general not severe enough. Of course, this is not to say that *all* physical or psychological threats are more coercive than any non-physical threats. The claim is, instead, that if a threat provides a full excuse, it is likely that the threat is physical or psychological.

It is possible, however, for an individual to be coerced into committing a wrongful act, not only as a result of a threat made by another person, but also through pressure exerted by incidental circumstances. As with intentional coercion, an agent subjected to incidental coercion must either commit an unjustified act or allow a harm to befall her.

And as with intentional coercion, incidental coercion provides an excuse for participating in a collective if 1) the agent will suffer serious bodily harm or death if she does not participate in the collective, 2) the serious bodily harm or death is otherwise unavoidable, and 3) the agent is not responsible for these circumstances. For example, if the only way for a starving individual to save her own life is to join a collective and consequently contribute to a collective harm, her primary responsibility for her contributions, and her secondary responsibility for the contributions of others might be mitigated as result of the overwhelming physical and psychological pressure caused by imminent starvation.

As with intentional coercion, there are many kinds of incidental coercion which might pressure an individual into joining a collective or to refrain from leaving it, but without imposing a risk of physical or psychological harm. As with non-physical intentional coercion, non-physical incidental coercion is unlikely to fully excuse the member, since there are few non-physical harms the imminent occurrence of which would overwhelm the will of an ordinary person. For example, the fact that the only way for an individual to prevent foreclosure on his house is to join a collective and consequently contribute to a collective harm would not fully excuse his primary or secondary responsibility.

7.7. PRIMARY AND SECONDARY RESPONSIBILITY COMBINED

Now that I have presented various factors that diminish primary and secondary responsibility, I will illustrate, by way of examples, how the criteria for secondary and primary responsibility combine to produce a principled account of individual

responsibility for collectively committed harms -- one which respects but does not follow blindly common sense intuitions. In the first section I focus on paradigmatic examples of membership in a collective. In the second section I focus on atypical examples. After this, I consider how responsibility -- especially secondary responsibility -- proliferates in nested collectives.

7.7.1. TYPICAL MEMBERSHIP

Consider again Pete, who is an employee of Evil Inc., a corporation that aims at generating profit by selling cigarettes to minors in under-developed countries. Pete's functional role is to solve computer-related problems in Evil Inc. Acting according to his functional role contributes to the corporation's ultimate goal of selling cigarettes.

The eligibility criteria of Pete's functional role are satisfiable by a large number of individuals, which means that declining to take the job will not significantly decrease the propensity of the collective act's occurrence. And Pete is one of many IT technicians who help troubleshoot email and network-related problems in the offices Evil Inc. Pete's modest contributions are superfluously overdetermining, which means that his participation only slightly increases the propensity of the collective act to which he contributes. And finally, Pete's contributions are highly subsidiary with respect to the ultimately object of his functional role. He is causally far removed from the profitable sale of cigarettes to minors in underdeveloped countries. The upshot is that Pete bears little primary responsibility for the object of his functional role, since the eligibility

criteria of his functional role are relatively easy to satisfy, and because his contributions are superfluously overdetermining and highly subsidiary.

But suppose Pete has a desire that Evil Inc. successfully maximize its profits by selling cigarettes to minors in underdeveloped countries. Though Pete has no particular desire for these minors to suffer, he values the profitability of Evil Inc. over their well-being. As a result, he has a desire that Evil Inc. successfully maximize its profits, even if this comes at the cost of the well-being of those to whom the cigarettes are sold. Since Pete is a low-level employee in the corporation, he does not have the power to determine whether Evil Inc. achieves its goals. But he contributes to these goals -- i.e., the object of his functional role -- in order to increase the probability that Evil Inc. achieves the goals of maximizing its profits.

But this is not the only reason that Pete intentionally contributes to the object of his functional role. Pete takes the job as an IT technician also because it pays well. So contributing to the object of his functional role is both a means to getting paid, and a means of increasing the probability that Evil Inc. will achieve its goals. Thus Pete has two instrumental reasons to contribute to the aim of selling cigarettes to minors.

That Pete 1) is a willing member of Evil Inc., in that he agrees to take on a functional role, 2) knows that he is contributing to the aim of selling cigarettes to minors, 3) intentionally contributes to this aim, and 4) desires the achievement of this aim, results in maximum secondary responsibility for what others do in furtherance of this aim, even

though Pete's functional role is relatively unimportant. As I have argued, voluntarily taking on a functional role is a basis for secondary responsibility for the contributions made by other members. According to the Functional Role Argument (see section 4.4.3), by making it his purpose to help further the shared object of his functional role, Pete bears secondary responsibility for the actions committed by every other member *qua* member. This is because he has implicitly committed himself to helping and to accepting help from every other member of Evil Inc., with respect to their contributions. So Pete is responsible for the contributions that others make to the collective act of selling cigarettes to minors (as long as their contributions are in accordance with their functional roles). Pete bears the *maximum* degree of secondary responsibility for the contributions of others in the collective, because he intentionally and knowingly contributes to the aim of selling cigarettes to minors in order to increase the probability of achieving this aim, and because he desires the occurrence of the collective act.

Because Pete is fully responsible for all the other contributions to the collective act, the responsibility that Pete bears is tantamount to the responsibility for collective's actions *tout court*. If Evil Inc. sells 10 million cigarettes to minors in the last fiscal year, then Pete bears secondary responsibility for selling 10 million cigarettes to minors during the last fiscal year. This is in spite of the fact that the contribution Pete makes to the object of his functional role is superfluously overdetermining, and that the eligibility criteria of his functional role are not very restrictive.

But suppose Pete takes the job as an IT technician only because it pays well -- not because he has any particular desire to contribute to the maximization of the corporation's profits. Acting according to his functional role is a means to getting paid. Thus Pete has an instrumental reason to act according to his functional role. As discussed in section 4.4, intentions are not closed under causal implication. That is, the fact that Pete recognizes he has an instrumental reason to act according to his functional role, and the fact that he intends to act according to his functional role, does not mean that Pete has an intention to contribute to the object of his functional role -- even if Pete is aware that acting according to his functional role results or is likely to result in a contribution to the corporation's aim of maximizing profits by selling cigarettes to minors. Even if the object of his functional role is not salient to Pete, he nonetheless commits himself to contributing to the object of his functional role, as a result of the Contributory Commitment Principles (see section 4.4.2). By committing himself to contributing to this object, Pete bears secondary responsibility for the contributions made by the other members, due to the Functional Role Argument (see section 4.4.3).

In such a case, Pete does *not* bear full secondary responsibility for the contributions made by all the other members. This is because he has no particular intention, instrumental or otherwise, to contribute to the success of the corporation's goal of maximizing profits by selling cigarettes to minors. And he has no particular desire, instrumental or otherwise, for the success of the corporation's goal. His lack of the relevant desire and intention does not prevent him from bearing secondary responsibility -- but it does prevent him from bearing full secondary responsibility.

But suppose that Evil Inc. has a policy of closely monitoring employees, to ensure that they are indeed contributing to the collective. If a monitoring supervisor discovers that an employee is not contributing to the collective, that employee will be fired, even if he is acting in accordance with his functional role. In this case, Pete, since he has a motivating reason to keep his job, has an instrumental reason not only to act according to his functional role, but to contribute to the object of his functional role as well. Pete is instrumentally rational -- he thus adopts an intention to contribute to the object of his functional role, even though he has no particular desire for the achievement of the corporation's aim.

As a result, Pete bears more secondary responsibility in this case than he did in the previous case. This is because, in the latter case, Pete intentionally contributes to the object of his functional role. He doesn't merely *implicitly* accept and provide help from other members -- he does so intentionally. He thus bears greater secondary responsibility than he would if he did not contribute intentionally. But like the previous case, Pete has no particular desire, instrumental or otherwise, for the success of the corporation's goal of maximizing profits by selling cigarettes to minors. As a result, he does not bear full secondary responsibility.

In the three cases discussed so far, Pete knows that acting according to his functional role is likely to contribute to the aim of selling cigarettes to minors. Pete bears even less secondary responsibility for the contributions made by others in the collective, if he is

culpably uncertain whether the corporation is selling cigarettes to minors. Pete's ignorance of whether acting according to his functional role will result in a contribution to a harm, might be either object-based or contribution-based. But either way, if he is aware of such a risk, i.e., that by acting according to his functional role he might contribute to the sale of cigarettes to minors, then to nonetheless act according to his functional role even though he has no particular wish to contribute to the aim of selling cigarettes to minors, is to commit a reckless contributory act. Again, Pete implicitly commits himself to contributing to the object of his functional role by taking on that role. By so committing himself to contributing to the aim of the collective, Pete bears secondary responsibility for the contributions made by the other members. But since his contribution is reckless, the secondary responsibility he bears is even less than what he would bear if he *knowingly* provided and accepted help in furtherance of the shared object of his functional role.

Consider a case in which Pete is culpably unaware of either the fact that the corporation is selling cigarettes to minors, or of the fact that by acting according to his functional role he contributes to the sale of cigarettes to minors. If Pete's culpable ignorance is contribution-based, then a reasonable person would recognize that acting according to such a functional role involves a substantial risk of a successful contribution to its object. After all, it is Pete's *purpose* to contribute to the object of his functional role. And while Pete is aware that he is acting according to his functional role, and that his purpose in so doing is to contribute to the sale of cigarettes to minors, he mistakenly and culpably believes that his role is so subsidiary, that he does not run a substantial risk of

contributing to the object of his role. As a result, if Pete actually contributes to the object of his functional role, the contribution is negligent, where the object of his culpable ignorance is the likely causal link between acting as part of the collective, and contributing to the harmful collective act. If Pete's culpable ignorance is object-based, then Pete unreasonably and mistakenly believes that the object to which he contributes is not the sale of cigarettes to minors. As a result, he unknowingly risks contributing to an unjustified collectively committed harm, even though he has no particular desire to contribute to such a harm. But whether Pete's ignorance is contribution-based or object-based, his ignorance of this risk is culpable, since a reasonable person would be aware of it. The resulting contribution to the object of his role is a *negligent* contribution.

Since Pete implicitly commits himself to contributing to the object of his functional role by taking on that role, he bears secondary responsibility for the contributions made by the other members. But since his contribution is negligent, the secondary responsibility he bears is even less than what he would bear if he *recklessly* provided and accepted help in furtherance of the shared object of his functional role.

I do not attempt to provide an absolute, (i.e., non-comparative) assessment of how much primary and secondary responsibility Pete bears in each of the above scenarios. The goal in this section is not to provide a precise metric of primary and secondary responsibility, but instead to present examples demonstrating how factors such as the causal relation between a contribution and its object, as well as pertinent intentions and beliefs, can affect primary and secondary responsibility.

7.7.2. ATYPICAL MEMBERSHIP

So far I have discussed paradigmatic cases of individual membership in collectives. In this section I will discuss how primary and secondary responsibility combine in atypical cases. Specifically, I will focus on cases in which membership in a harmful collective contributes nothing to that harm, and has good consequences for the member.

Suppose ψ and Σ are both events, where the occurrence of Σ is bad and the occurrence of ψ is good. Suppose that, when the two are weighed, the bad which Σ causes outweighs the good which ψ causes. k is a member of a collective that commits Σ . k 's role in the collective is to contribute to the occurrence of Σ by committing ϕ . And ϕ , in addition to contributing Σ , also causes ψ . So by contributing to Σ , k causes a good event, ψ .

Suppose that the only way for k to contribute to Σ is to commit ϕ . On strictly consequentialist grounds, the fact that Σ causes a worse state of affairs than ψ does not imply that it is impermissible to *contribute* to Σ . For a consequentialist, the bad resulting from *contributing* to Σ must be weighed against the good resulting from causing ψ .

Suppose, for example, that k works at a corporation that specializes in manufacturing chemical weapons. The aim of this corporation is to maximize its profit by selling chemical weapons to the government, which has a history of unjustifiably using chemical weapons during wars. By committing the set of tasks required of her as an employee -- i.e., by committing ϕ -- k contributes to the achievement of the corporations aims, Σ . But

by committing ϕ , k also earns an income which financially sustains k and her family -- i.e., by committing ϕ she causes ψ . We can assume that for the corporation to meet its goals will result in a far worse state of affairs than would be the case if k were forced to find work elsewhere -- or even if k were forced into unemployment. That is, the occurrence of Σ is far worse than the non-occurrence of ψ .

To determine the permissibility of contributing to Σ by committing ϕ , a consequentialist would wish to know whether Σ would occur if k did not commit ϕ . That is, a consequentialist would wish to know whether quitting her job would make a difference with respect to the occurrence of Σ . But on my view, when an agent acts as a member of a collective, she is responsible not only for the difference that she makes, but also -- to a degree depending on details such as her relevant mental states -- for the difference made by the other members. As a result, k is responsible for potentially far more than the difference that she makes. Given this secondary responsibility, and given the badness of Σ , it is unlikely that membership in the collective is justified, even though such membership involves committing ϕ which furthers the well-being of k 's family.

Suppose that though k acts according to her functional role, she does not causally contribute anything to Σ because her employers did not competently draft the goal-oriented rules constitutive of the functional role. Even in such a case, the secondary responsibility she bears is far greater than the degree that she would bear if we considered only the difference that her contribution makes to Σ . Since her responsibility for the contributions of others is largely the result of agreeing to act according to a functional

role in the collective, she need not contribute anything in order to bear secondary responsibility.

The upshot is that it is impermissible for k not only to contribute to Σ as a member of the chemical weapons corporation, but also to be an employee of the corporation, even if the difference she makes is little or none, and even if non-membership has bad consequences for her and her family. However, if k , upon joining the collective, has no desire for the occurrence of Σ , and if she has no particular intention to contribute to the object of her functional role, then she bears less secondary responsibility for the contributions of others in comparison to those who do indeed desire the occurrence of Σ , and those who do not desire Σ but who intentionally contribute to it. Recall from section 7.4 that unshared contributory intentions result in less secondary responsibility than shared contributory intention. Nonetheless, joining the collective is likely to remain morally wrong, even though the alternative is detrimental to k 's family.

But what if k intentionally contributes to both the achievement and to the failure of Σ ? Suppose k both acts according to her functional role and sabotages the corporation's efforts by committing an act that reduces the tonnage of chemical weapons manufactured this quarter. We can characterize the resulting event as the occurrence of Σ^* , which is a bad event, but not as bad as Σ . k bears secondary responsibility for the contributions made by others in the collective. However, k also bears *primary* responsibility for the moral difference between Σ and Σ^* . Since k bears some secondary responsibility for the contributions made by others in the collective, it seems that k bears significant

responsibility for the harm done by Σ^* , even though she prevented Σ . If k bears full secondary responsibility for the contributions made by others, then k's primary and secondary responsibility combined is the positive effect of bringing about Σ^* instead of Σ , summed with the negative effect of Σ^* over $\sim\Sigma$. If the positive difference between Σ^* and Σ is greater than the positive difference between Σ^* and $\sim\Sigma$, then k's positive contribution outweighs the secondary responsibility resulting from acting according to her functional role. But if the positive difference between Σ^* and Σ is *less* than the positive difference between Σ^* and $\sim\Sigma$, then k's positive contribution is *outweighed by* her secondary responsibility resulting from acting according to her functional role.

This is intuitively implausible. It seems that even if k reduces the total output of chemical weapons by just a little -- and contributes nothing at all to the production of the chemical weapons -- that k is not responsible for the chemical weapons that are manufactured. But on my account, it seems that for k to avoid net positive responsibility for the manufacturing of chemical weapons, she must reduce the total output by an amount that does more good than would be done if the new amount of chemical weapons produced were brought to zero. This is a very strong requirement. But since, for any amount of chemical weapons the corporation produces, k wants to *reduce* that amount, k does not desire the achievement of the corporation's aims. In addition, k does not share a desire with anyone for the achievement of the corporation's aims. Nor does she intend to provide a *net* contribution to the object of her functional role. This diminishes her secondary responsibility for the contributions made by others since she does not share contributory intentions with any other members of the collective (see section 7.4). As a

result, even if the positive difference between Σ^* and Σ is *less* than the positive difference between Σ^* and $\sim\Sigma$, k 's positive contribution might nonetheless outweighs her secondary responsibility resulting from acting according to her functional role.

But there is another reason to believe that even if the positive difference between Σ^* and Σ is less than the positive difference between Σ^* and $\sim\Sigma$, k does not bear secondary responsibility for the contributions made by others in the collective. Put simply, if k is intentionally not contributing to the collective, and if she is actively sabotaging the contributions made by others in the collective, then there might be grounds to claim that she is not an actual member of that collective. She is instead a pseudo-member. After all, she not only flouts her functional role, but actively works against its object as well. If this is the case - if she is not an actual member - then she does not bear secondary responsibility for the contribution to Σ made by others -- but she retains primary responsibility for her efforts at sabotaging Σ .

7.7.3. RESPONSIBILITY IN NESTED COLLECTIVES

Many collectives consist of smaller collectives. It is worth noting how responsibility proliferates in such collectives. A large collective consisting of smaller collectives is a *nested* collective. The constituent collectives are sub-collectives, and the constituted collective is the super-collective. As I will discuss in section 8, a nation's military is an obvious instance of a nested collective. In this section, I will consider how primary but especially secondary responsibility proliferates in nested collectives.

But first, it is necessary to further elucidate the structure of nested collectives. Of the n sub-collectives composing a nested collective, it need not be that $n-1$ of the sub-collectives are themselves nested. That is, the n collectives might be arranged like a jigsaw puzzle, rather than like Matryoshka dolls. Of course, a nested collective might be a hybrid of the two pure forms.

Each sub-collective constitutive of a larger collective can be regarded as a collective in its own right; it has its own Σ and its own set of functional roles, and thus the members of the sub-collective bear secondary responsibility for the contributions made by one another to the shared object of their functional roles. However, the functional roles constitutive of the sub-collective are by design logically concomitant with the functional roles constitutive of the super-collective. As a result, having the former functional role logically entails having the latter. Moreover, acting according to the former functional role logically entails acting according to the latter functional role. This structural feature reflects the fact that members of the sub-collective have purposes both within that sub-collective and within the super-collective.

Suppose a collective, C , consists partly of sub-collective B , which consists partly of sub-collective A . Consider the following two ways that a member of B might bear secondary responsibility for what the other sub-collectives do. First, the member might bear responsibility for the act committed by the sub-collective which B subsumes -- viz., A . Second, the member of B might bear responsibility for the act committed by the sub-

collective that subsumes B -- viz., C. For convenience, call the first type of responsibility 'downstream responsibility', and the second type, 'upstream responsibility'.

There are intuitive reasons to believe that members of a sub-collective bear both kinds of responsibility. In a basic chain of command, higher ranked individuals are typically responsible for the actions of lower ranked members. In a collective organized according to a nested chain of command, subordinate members occupy sub-collectives that are nested relative to the sub-collectives occupied by their superiors. Suppose A, B, and C constitute a nested chain of command, in which A commands B and B commands C. All things being equal, the degree of responsibility that members of B bear for what members of C do, is greater than the degree of responsibility that members of B bear for what members of A do. In such a collective, there is *downstream* responsibility.

However, there is a straightforward sense in which members of a sub-collective such as B bear *less* responsibility for the acts committed by the sub-collective B subsumes -- A -- than the sub-collective B is subsumed by -- C -- since some of B's members are *not* members of A, but *all* of B's members are members of C. If the sub-collective A commits a wrongful act, only *some* of B's members were in the sub-collective that contributed that act. But if the collective that subsumes B commits a wrongful act, then *all* of B's members were in the sub-collective that committed that act. For this reason, B seems to bear greater responsibility for an act committed by the sub-collective that subsumes B, than for the act committed by the sub-collective which B subsumes. Here, responsibility is *upstream*.

Upstream and downstream responsibility are not incompatible -- one and the same members of a sub-collective can bear both upstream and downstream responsibility. A basis for downstream responsibility is the causal influence a collective has on the sub-collective it subsumes. A nested collective organized according to a chain of command is an example of this. A basis for upstream responsibility is the mereological relation between a sub-collective and a higher-order collective. I will argue that secondary responsibility accounts for upstream responsibility, and primary responsibility accounts for downstream responsibility.

Suppose collective RA consists of members m1 through m9. RA subsumes a collective RB which consists of members m4 through m9. RB, in turn, subsumes RC, which consists of members m7 through m9. The ultimate shared object of the functional roles in the super-collective RA, is $\Sigma 1$. Suppose RC is charged with contributing to $\Sigma 1$ by committing $\Sigma 3$, which promotes $\Sigma 2$. Likewise, RB is charged with contributing to $\Sigma 1$ by committing $\Sigma 2$, which promotes $\Sigma 1$. This collective can be described as follows:

RA(m1, m2, m3, RB(m4, m5, m6, RC(m7, m8, m9))) commits $\Sigma 1$

RB(m4, m5, m6, RC(m7, m8, m9)) commits $\Sigma 2$

RC(m7, m8, m9) commits $\Sigma 3$

All members of RA bear secondary responsibility for $\Sigma 1$, $\Sigma 2$, and $\Sigma 3$. If a member of RC acting according to her functional role makes a contribution to $\Sigma 2$, she *eo ipso* contributes

to Σ_1 , which is the ultimate object of her functional role. Members of RB are also members of RA. Members of RA bear secondary responsibility for contribution to Σ_1 , Σ_2 , and Σ_3 . Therefore, members of RB, qua members of RA, also bear secondary responsibility for contribution to Σ_1 , Σ_2 , and Σ_3 . But members of RB, qua members of RB, share secondary responsibility for contributions to Σ_2 . Since a member of the sub-collective acts both qua member of the sub-collective and qua member of the super-collective, the member bears secondary responsibility for the acts committed by the super-collective and responsibility for the acts committed by the sub-collective. Secondary responsibility resulting from these two sources is *additive*.

If an individual is a member of the super-collective, then that individual bears some secondary responsibility for all that the super-collective does -- and part of what the super-collective does is what the sub-collective does. If, in addition, the individual is a member of the sub-collective constitutive of the super-collective, then she bears some secondary responsibility, once again, for what the sub-collective does. What the sub-collective RB does, in the current example, is to commit Σ_2 . As a result, members of RB bear secondary responsibility *twice over* for Σ_2 , in comparison to members of RA who are not members of RB. Members of RC are also members of RA and RB. So members of RC bear responsibility for what RA does, what RB does, and, of course, what RC does. So members of RC, by acting qua members of RA, bear secondary responsibility for Σ_1 , Σ_2 , and Σ_3 . Members of RC, by acting qua members of RB, they bear secondary responsibility for Σ_2 . And by acting qua members of RC, bear secondary responsibility

for $\Sigma 3$. So members of RC bear secondary responsibility for $\Sigma 1$, and they bear secondary responsibility *twice over* for $\Sigma 2$ and $\Sigma 3$.

But what does it mean to be responsible ‘twice over’? The responsibility at issue here is obviously *secondary* responsibility. An individual who bears secondary responsibility for an event might be liable to provide compensation for harms caused by that event. The amount or degree of compensation that the individual is liable to provide depends, of course, on the harm for which the liable party is responsible; but it also depends on the degree to which the liable party bears responsibility for that harm. If the liable party bears very little responsibility for the harm, then the liable party will, generally, have to provide very little compensation for that harm. If the liable party bears twice the responsibility for that same harm, then the liable party will, generally, have to provide greater compensation for that harm. It is perhaps tempting to claim that because the liable party is twice as responsible, then the amount which the individual must pay in compensation is twice as great. But there is no decisive reason to believe that responsibility and liability to compensate are linearly related in this way.

For a member to be responsible twice over for a harm also has consequences for liability to preventive injury. When a large nested collective is committing a harmful act, it is unlikely that every member who bears some responsibility for the ensuing harm will be in a position to prevent that harm. If a member *were* in such a position, in that harming the member would prevent the collective harm, then the amount or degree of preventive harm

to which the member is liable depends on the degree of responsibility that the member bears for that harm. Obviously, it is difficult to quantify not only responsibility but harms as well -- which makes it difficult to determine precisely how much preventive harm a responsible individual is liable to suffer. But it is safe to say that a member who bears twice the amount of responsibility for a collectively committed harm than he would otherwise bear is liable to a greater amount of preventive injury than to which he would otherwise be liable.

In addition, a party responsible for an impermissible harm is typically thought to be punishable for that harm. As mentioned in section 2, if a member bears a small degree of secondary responsibility for a harm committed by a collective, and if that responsibility is a basis for blaming the individual for the harm, and if that blame is a basis for punishment, then the individual will be liable to a small punishment. If the individual bears twice the amount of responsibility, then the individual will be liable to a *larger* punishment (though, not necessarily a large punishment). In any case, if the degree of secondary responsibility that a member bears is small, then bearing that responsibility twice over will make little difference.

In any case, it might seem that I have argued that those members who occupy the bottom rungs of a hierarchical collective bear the *most* responsibility even though they intuitively bear the *least* responsibility. This is ludicrous result. But fortunately, this is not a consequence of my view. While a member of RA would indeed bear greater secondary responsibility than someone higher up in the chain of command, she would bear

significantly less *primary* responsibility. Those who occupy sub-collectives that rank high in a chain of command, such as members of RC, bear significant primary responsibility for the contributions made by those below them, due to the tremendous causal and counterfactual significance of their contributions, viz., the occurrence of $\Sigma 1$. If the lower ranking sub-collectives, as a whole, have little causal influence on $\Sigma 1$, then secondary responsibility for their contributions will be accordingly minimal -- and even this degree of secondary responsibility twice over will not be nearly as great as the degree of primary responsibility that more influential, higher ranking collectives will bear.

Put differently, upstream and downstream responsibility not only have different valences, but have different magnitudes as well. In general, downstream responsibility is stronger. But upstream responsibility is still crucial in that it reveals the influence that nested compositionality has on secondary responsibility. Secondary responsibility, because it has a non-causal basis, has a tendency to 'spread out' evenly in a collective. This is not so, however, in a nested collective; the dissemination of secondary responsibility is sensitive to the compositional structure of such collectives in that there is greater upstream than downstream secondary responsibility.

Together, upstream and downstream responsibility undergird the intuition that members of a sub-collective bear the greatest responsibility for what that sub-collective does -- more so than for what other sub-collectives do. Consider a typical nested collective -- members of sub-collective R_i will bear secondary responsibility twice over for what that sub-collective does and for what successive subsumed sub-collectives do. Those

members will not bear secondary responsibility twice over for sub-collectives subsidiary to R_i . Since primary responsibility depends largely on the causal and counterfactual effects of a contribution, and since such effects tends to dissipate as the contribution causally propagates, it stands to reason that, all thing being equal, primary responsibility among members of R_i will be greatest for what those members do, and will diminish with successive subsidiary sub-collectives. As a result, for any sub-collective, there is a ‘nexus’ of responsibility, in that the combination of primary and secondary responsibility will be greatest for the acts committed by fellow sub-collective members, and will diminish upstream and downstream. This result accords with the powerful intuition that individuals in a sub-collective bear greater responsibility than those outside it for what that sub-collective does.

Of course, it might be argued that primary responsibility alone can undergird this intuition. But for reasons I presented in section 6.1, an appeal to causal influence is not enough to ground the intuitive responsibility that members of a collective bear for what they together do. Secondary responsibility is not only compatible with but also undergirds the intuition that small groups in large collectives bear greater responsibility for what *they* do than for what other individuals in the collective do.

As a final side-note on how nested compositionality affects responsibility, recall also that intentions also affect the degree of secondary responsibility that a member bears. A member of RB is in a position to *intentionally* act as part of RB . This is in contrast to a member of RA who is not a member of RB . If the member of RB does indeed

intentionally act as part of RB, then she bears greater secondary responsibility for the actions committed by $\Sigma 2$ than the member of RA does. The absence of such an intention, as I argued in section 4.4, does not preclude membership in the pertinent collective; nor does it preclude, as I argued in section 7.5, secondary responsibility for the contributions made by others who participate in the collective. However, an individual who has an intention to contribute to the object of a functional role bears greater secondary responsibility, relative to the degree the individual would bear if she did have such an intention.

8. War - a Test Case

In this section I apply the account of individual responsibility for collective action to a paradigm type of cooperative collective action: war. Specifically, I will be concerned with the responsibility of individual combatants for acts committed by the military of a country waging a war, where these acts are defined broadly enough to include the decisions made by military and civilian leaders. Since the military tends to encompass multiple, complex collectives, I will begin by elucidating the organizational relationships constitutive of modern militaries.

An intermediate conclusion will be that a just war cannot involve the pursuit of unjust subsidiary aims by the military or by any of its constituent sub-collectives. Ultimately, by applying the account of individual responsibility for collective action, I hope to show that participating in a just element of a war can be impermissible. That an element of a war is just does not mean that participating in that element is permissible. Likewise, participating in an unjust element of a war can be permissible. That an element of a war is unjust does not mean that participating in that element is impermissible. Unsurprisingly, if we ignore secondary responsibility, combatants bear less responsibility for the conduct of other combatants.

8.1. THE STRUCTURE OF WAR

Here I address the conceptual relationship between aims and wars. Following the practice of just war theorists, my use of the term 'war' will distinguish between the sides in a

conflict. For example, what we call ‘the Franco-Prussian War’ consisted of France’s war against Prussia and Prussia’s war against France. These were, in one sense, distinct wars. This is the sense of ‘war’ that I will use here.

Waging a war involves the pursuit of aims through the application of military force. I will make a conceptual distinction between various types of aims. The *ultimate* aims of a war are those that explain the resort to war. These aims are the motivating reasons for the government’s resort to war. When referring to the aims of a war, I will assume that the aims are ultimate, unless stated otherwise.

An ultimate aim subsumes subsidiary aims. These are aims the achievement of which is intended to constitute or cause the achievement of particular ultimate aims. Subsidiary aims can be broken down further, into operational objectives which subsume tactical objectives, which, in turn, subsume missions. For example, an ultimate aim of a war might be to annex land rich in oil. An operational objective might be to destroy enemy airbases bordering the area to be annexed. A tactical objective might be the destruction of hardened targets in a particular enemy airbase. A mission might consist in a particular air strike against the hardened targets using precision-guided munitions.

Several additional features of subsidiary and ultimate aims will be especially relevant. First, those planning the war might reason speciously when choosing subsidiary aims. And even instrumentally justified subsidiary aims can fail to yield the appropriate ultimate aim due to unforeseeable circumstances. But what makes an aim subsidiary is

not its effectiveness in yielding an ultimate aim. Rather, for an aim to be subsidiary, it is enough that it is *intended* as a means to an ultimate aim.

Second, the identity of an ultimate or subsidiary aim does not necessarily depend on the aims that are subsidiary to it. So if a particular aim, s_1 subsumes an aim subsidiary to it, s_2 , we can substitute s_2 with other possible subsidiary aim, without necessarily changing s_1 .

Third, for any war, given a government's particular set of ultimate aims, there are various *other* sets of ultimate aims that the government could have chosen instead of the set it actually chose. Some possible sets of aims have so little in common with one another that it is more appropriate to regard the wars they compose as different wars altogether rather than the same war with different aims. But in general the same war can have different aims, within limits. In these cases, the government can alter the war's aims without necessarily waging a different war.

Fourth, ultimate and subsidiary aims are not necessarily exclusive. It is possible for an aim to be both ultimate and subsidiary. In such a case, the aim is both a reason for waging war and an intended means to the achievement of some other ultimate aim.

It is, of course, an idealization to claim that governments adopt a particular set of aims for a war. There are varying degrees of commitment towards the pursuit of particular aims, and leaders are often capricious in their commitments. Moreover, the aims adopted might

be indeterminate or ill-defined. The aims of a war are usually the result of collective decision-making; this can introduce indeterminacy with respect to the war's aims, depending on the decision-making procedure that the collective uses. Sometimes the aims are intentionally left vague in order to reach consensus among members of a gridlocked government, or to facilitate post-bellum claims of success; sometimes they are unintentionally vague simply as a result of unreflective leadership. I believe the arguments I will provide can be amended to fit these scenarios. But for the sake of simplicity, the hypothetical wars that I will discuss will have clear and stable sets of ultimate aims.

8.2. TYPES OF UNJUST AIMS

So far I have discussed only the structural relationships among aims in a war. Now I turn to the moral evaluation of those aims. The following claims will be highly generalized, so that my ultimate conclusions will be compatible with a variety of theories of jus ad bellum (i.e., the conditions according to which a resort to war is justified).

There are, broadly construed, three reasons why pursuing an aim can be impermissible. First, pursuing an aim might be unjust, and therefore impermissible, if the aim itself involves the violation of rights. In such cases there are typically no methods of achieving the aim permissibly. *A fortiori*, military force is generally impermissible as a means to achieving such an aim. Call such aims 'intrinsically unjust'. Genocide is an obvious example of an intrinsically unjust aim.

Second, pursuing an aim can be impermissible because it violates constraints of proportionality. Jeff McMahan argues that there are two kinds of proportionality-violations. Sometimes an agent commits a harm or wrong for which that agent is liable to be harmed. But to kill the agent might be excessive or disproportionate given the basis of her liability. Hence the agent is not liable to be killed. McMahan calls this a constraint of ‘narrow proportionality’. An example of an aim the pursuit of which would violate narrow proportionality is that of marginally improving the status of women, for which no one bears enough liability to be justifiably killed.³

Alternatively, an aim might violate what McMahan calls a constraint of ‘wide proportionality’, in which the good of the aim being pursued (and perhaps the good side-effects of its pursuit) is weighed against the harms caused to innocent or non-liable people, usually as side-effects of pursuing the just aim (2007). For example, collateral damage to civilians during a tactical bombing of a munitions factory must be weighed against the good of destroying the munitions factory for the bombing to satisfy constraints of wide proportionality.

There is another way in which the pursuit of an aim by a particular means is impermissible. A particular means to the accomplishment of an aim can satisfy constraints of wide and narrow proportionality, and still be impermissible to pursue, if

³ This example belongs to Thomas Hurka (2002). Its explication in terms of liability belongs to Jeff McMahan.

there is an even less harmful means to the accomplishment of the same aim, or some other aim that would make an equal contribution to the achievement of the just cause. Satisfying proportionality might not require pursuing the *least* harmful means. But a failure to pursue the least harmful means might still be wrongful. If so, then there is a constraint on the means to the pursuit of an aim independent of the constraints of proportionality, viz. that it must be the least harmful means available.

To summarize, I have distinguished three ways in which pursuing an aim can be impermissible. An aim might be intrinsically unjust. Or the particular means of pursuing an aim might violate constraints of (wide or narrow) proportionality. Or the particular means might be unnecessarily harmful. Having drawn these distinctions, I will elucidate the relationship between the moral status of a war's aims, and the moral status of the war itself.

I will consider individual responsibility for participation in several varieties of wars. I will, of course, consider just wars and unjust wars. But I will also distinguish between homogenous and heterogeneous unjust wars. The former is a just war consisting solely of unjust aims and subsidiary elements. A heterogeneous unjust war consists of both just and unjust aims. Among heterogeneous unjust wars, I further distinguish between *narrowly* unjust and *broadly* unjust wars. The former but not the latter is an unjust war for which it is foreseeable that it will have positive consequences that outweigh its negative consequences.

8.3. THE MILITARY AS A NESTED COLLECTIVE

Primary responsibility in homogenous wars will depend chiefly on the importance of the combatant's functional role, the consequences of the combatant's contributions to the war, as well as the combatants mental states, i.e., his intentions and beliefs. Put very basically, a combatant bears primary responsibility for events that would not have occurred had she not been a member. If the combatant is a ranking officer, then she can be characterized as a member of a sub-collective in the greater collective consisting of the military as a whole. She bears primary responsibility for the acts committed by those under her command; they can be characterized as members of a sub-collective subordinate to their commanding officer's sub-collective. Put differently, the commanding officer will bear what I have called 'downstream' responsibility. Since lower-ranking officers and non-commissioned combatants typically have less of an influence over the actions of their commanding officers than the commanding officers have over them, the former will bear less downstream responsibility than the latter do. And since lower ranking and non-commissioned combatants are more easily replaceable than higher-ranking officers, and because the contributions they make are often superfluously overdetermining, these combatants will bear little primary responsibility. But these combatants -- as well as ranking officers -- will bear secondary responsibility for the acts committed by the rest of the military, regardless of how much they contribute. The degree to which a combatant bears secondary responsibility will depend partly on her mental states, such as her intentions. They will also bear secondary responsibility twice

over for the acts committed by those in the sub-collective of which they are a part -- that is, they bear what I have called 'upstream responsibility'.

The specific events for which any given low-ranking or non-commissioned combatant in a war will be responsible will depend on that combatant's particular circumstances. We can, at best, make generalizations about the magnitude of responsibility that any such combatant bears, and the causal reach of different types of combatants. Since every ultimate and subsidiary aim of a homogeneously unjust war is itself unjust, it is likely that most combatants will bear some primary responsibility for the wrongs committed by the collective of which he is a part. As a result, most combatants will bear limited primary responsibility for the wrongs in virtue of which a war is unjust -- and this is to say nothing of the exculpating influence of non-culpable ignorance, coercion, etc.

Secondary responsibility, however is not limited in this regard. The degree of secondary responsibility that a combatant bears in a morally homogenous war depends in part on the organizational features of the military in which the combatant serves, and the position in the military that the combatant occupies. There are many mereologically respectable ways to categorize the elements of the armed forces. But the most well-known and useful categorization consists of a hierarchically organized multiply-nested collective. The sub-collective that a combatant occupies will affect the combatant's secondary responsibility in general. Secondary responsibility in a nested collective is stronger upstream than downstream -- the reverse of primary responsibility. For example, in the US, the army is organized as follows:

type	persons	constituent units
region, theater, or front	1,000,000+	4+ army groups
army group	250,000+	2+ armies
army	60,000–100,000+	2–4 corps
corps	30,000–80,000	2+ divisions
division	10,000–20,000	2–4 brigades or regiments
brigade	2000–5000	2+ regiments, 3–6 battalions
regiment or group	2000–3000	2+ battalions
infantry battalion,	300–1000	2–6 companies, batteries
infantry company, artillery battery	70–250	2–8 platoons
platoon	25–60	2+ squads or sections
section or patrol	8–12	2+ fireteams
squad or crew	8–16	2+ fireteams or 1+ cell
fireteam	4–5	n/a

Suppose *c* is a member of a platoon. This platoon is a sub-collective; suppose it partly composes the sub-collective consisting in an infantry battalion. The 25 to 60 soldiers composing the platoon, provided they act according to their functional roles, act not only qua members of the army *tout court*, but qua members of the platoon, and qua members of the infantry battalion. Each of the 25 to 60 soldiers, as with any soldier in the armed forces, would bear secondary responsibility for the actions committed by other combatants, qua combatants. Because *c* is a member of a platoon, an infantry battalion, and an army, and since those who act qua members of the platoon *eo ipso* act qua members of the infantry battalion and the army, *c* bears secondary responsibility for the contributions of her fellow platoon members and for the contributions of her fellow battalion members twice over (see section 7.7.3).

Moreover, since the platoon is a collective unto itself, the members are in a position to intend to act not only as part of the armed forces of their nation, but of the platoon as well. Recall that intentions affect the degree of secondary responsibility that a member

bears. *c* is in a position to intend to act not only as part of the armed forces of their nation, but of the platoon as well, if she has such an intention, then she bears greater secondary responsibility for the actions committed by others in her platoon, relative to those members of the armed forces outside the platoon. The absence of such an intention, as I argued in section 4.4, does not preclude membership in the pertinent collective; nor does it preclude, as I argued in section 7.5, secondary responsibility for the contributions made by others who participate in the collective. However, for an individual to have an intention to contribute to the object of her functional role adds to secondary responsibility relative to the degree the individual would bear if she did not have such an intention. So if *c* has an intention to contribute to the aim of the platoon of which she is a member, her secondary responsibility for contributions made by other members of the platoon will be greater, *ceteris paribus*, than those made by combatants who are not members of the platoon. *Mutatis mutandis* for membership in other sub-collectives composing the army.

So if a platoon is tasked with accomplishing a particular mission, ψ , then every member of the platoon will bear secondary responsibility twice over for the contributions made by fellow members of the platoon. If a member intends to contribute to this goal, then the secondary responsibility she bears for ψ will be even greater. If accomplishing ψ requires a contribution by an infantry battalion, and if it so happens that the platoon of which *c* is a member is also partly constitutive of the infantry battalion tasked with ψ , then *c* will bear secondary responsibility twice over not only for the platoon's contribution to ψ , but for the battalion's contribution as well.

Membership of those who participate in a particular element of a war might cross-cut membership in discrete military units. Suppose two patrols are charged with accomplishing mission Φ . If c is a member of one of these patrols, then she might be a member of at least two sub-collectives: the sub-collective consisting in the patrol of which she is a member, and the sub-collective consisting of those tasked with accomplishing Φ . In this case, c bears responsibility twice over for whatever contributions are made to the war by the members of her platoon, as well as those made by the sub-collective consisting of those tasked with accomplishing Φ . In this example, it is assumed that the two platoons are sufficiently independent of one another, in that they can be regarded as distinct sub-collectives composing the collective consisting of those tasked with committed Φ . This need not be the case, however. If the two platoons interact sufficiently, then the distinction between these token military units loses their pragmatic relevance. They can be regarded as a single sub-collective, consisting of those tasked with committing Φ .

There are, of course, other ways that collectives divided into military units might be related to collectives divided by elements of a war. In any case, a combatant belonging to any such collective will bear secondary, 'upstream' responsibility twice over for the acts of others in virtue of his membership in both the sub-collective and the army tout court. The combatant will bear primary, 'downstream' responsibility for the effects he has on the course of the war. The magnitude of primary responsibility resulting from directly causing an unjustified harm is much stronger than the magnitude of secondary

responsibility resulting from sharing membership in a collective a member of whom caused an unjustified harm qua member of the collective.

It is clear that the various sources of a combatant's responsibility will pull in different direction when the war in which the combatant participates consists of both just and unjust elements. In the next section, I consider individual responsibility in such wars.

8.4. NARROWLY VS. BROADLY UNJUST WARS

Most wars consist of both just and unjust elements. Even the most historically noble wars typically included *some* unjust element. Before I examine how a war with just and unjust elements affects the responsibility of individual combatants, I will examine the moral status of these wars themselves. Here I argue that a just war cannot include *any* ultimate or subsidiary aims that are unjust. I argue in favor of this view by way of example.

Suppose a government, as a result of civil unrest, embarks on a campaign of atrocities against its own population in order to deter further resistance. The government has its soldiers commit unspeakable acts against the civilian population indiscriminately. Call this country 'ATROCITY'. The government of a bordering country is considering military intervention in order to stop the massacres for humanitarian reasons. Call this country 'INTERVENE'.

The government of INTERVENE is considering launching a ground assault to neutralize the military units carrying out the massacres, most of which consist of the government's private guard. Suppose the government of INTERVENE is aware that this would effectively eliminate the means by which the country's despots maintain their control, allowing the people of ATROCITY, if they wish, to overthrow the government and replace it with a provisional one of their own choosing. The people of ATROCITY would welcome INTERVENE's assistance in stopping the massacres and the government of INTERVENE is aware of this. It is also aware that stopping the massacres will satisfy the constraints of necessity and (narrow and wide) proportionality.

However, INTERVENE is considering another aim, in addition to stopping the massacres. Bordering the two countries is a strip of unpopulated land belonging to ATROCITY. This borderland itself has little value but is strategically ideal for INTERVENE as a buffer between the two countries. Because of this, the government of INTERVENE is considering annexing this borderland, in addition to neutralizing the massacring military units. However, if the government of INTERVENE pursues both aims, it will seize the borderland first and only then stop the massacres.

Suppose that annexing the borderland will not directly harm any civilians. Pursuing the aim of stopping the massacres and the aim of annexing the borderland would yield a better outcome, impartially considered, than pursuing neither. Suppose further that the government of INTERVENE is aware that the people (though, of course, not the government) of ATROCITY would rather bear the violation of their rightful sovereignty

over the borderland than continue to be subjected to massacres by their government. This is not to say that the people of ATROCITY do not mind the annexation. It can be predicted that once the provisional government is in place, the people of ATROCITY will, via this government, protest the annexation. However, they will have neither the military, political, nor economic resources to reclaim the borderland.

The ultimate aim of annexing the borderland is intrinsically unjust. INTERVENE has no right to the territory. Acquiring it is not the sort of aim that can be pursued through military force, regardless of how few casualties are incurred. And the government of INTERVENE is, by hypothesis, in a position to adopt the aim of stopping the massacres *without* annexing the borderland. So annexing the borderland is not subsidiary to the just aim of stopping the massacres.

But suppose that because the government of INTERVENE would not benefit by pursuing only the aim of stopping the massacres, the government would rather do nothing, thereby allowing the massacres to occur, than go to war without pursuing the further aim of annexing the borderland. And suppose that it is permissible for INTERVENE not to go to war. Is this a reason to believe that a war with the aims of both stopping the massacres and annexing the borderland is permissible?

It is hard to see how this would be so, given that the aim of annexing territory is not subsidiary to stopping the massacres. The government of INTERVENE is, by hypothesis, free to pursue the aim of stopping the massacres without the aim of annexing the

borderland, if the government so wishes. Pursuing a set of aims that includes annexing the borderland is unjust partly *because* doing so involves freely rejecting an alternative set of aims that does not include the aim of annexing the borderland. A recalcitrant disregard for reasons not to commit a certain act typically does not diminish the reasons not to commit that act. So even if the government of INTERVENE will stop the massacres only if it also annexes the borderland, pursuing both aims is unjust. This is so despite the fact that pursuing both aims makes things better overall than they would be if the government of INTERVENE chose to pursue none of the aims.

A war that make things better overall relative to the absence of that war, yet is nonetheless unjust, can be called ‘narrowly unjust’. Unjust wars that do *not* make things better overall relative to the absence of that war, I will call ‘broadly unjust’. These classifications help reveal the moral heterogeneity of an unjust war’s aims.

It might be argued that a war is just if and only if going to war has better consequences than not going to war. On this view, a war resulting in an improvement over what would have been the case without that war is just. But this view is absurd. If a war is just if and only if it has better consequences than not going to war, then it is morally permissible for a government to ‘tack on’ gratuitously harmful, self-serving aims when waging otherwise just wars. On this view, if a country is the victim of unjust aggression, the government of that country can permissibly pursue aims that, for example, ignore duties of care, as long as pursuing this aim in combination with pursuing the aim of self-defense has better consequences than not going to war at all. Or if the government of a country (such as

INTERVENE) is waging a war with a humanitarian aim, the government can permissibly pursue aims wronging the people requiring assistance, as long as these wrongs are, for its victims, a small price to pay in comparison to losing assistance from the intervening power. But this is not just; this is extortion.

Of course, if achieving the aim of stopping the massacres is costly then INTERVENE might be entitled to compensation. For INTERVENE to be entitled to the borderland, the people of ATROCITY would have to agree to give it up to INTERVENE as compensation for military assistance. It is unrealistic to presume that the victims of an oppressive regime would have the political voice necessary to explicitly contract with a foreign power. As a result, perhaps it is permissible for INTERVENE to act according to a *hypothetical* contract; its terms are determined partly by what the people of ATROCITY *would* agree to, if they were to explicitly partake in such a contract. I will assume that the people of ATROCITY would indeed consent to sacrificing the borderland as a price for intervention. This does not mean, however, that INTERVENE is entitled to the borderland. Individuals often agree to contracts under duress or in conditions of extremity. Seeking agreement to a contract under such circumstances can be extortionate, if the price of the service offered is either excessive in relation to the cost of providing it or in relation to the value of the service itself. The same might be said of hypothetical contracts in which the hypothetical agreement is made under conditions of duress. We can assume both that the people of ATROCITY are under duress, and that the strategic value of the borderland is significantly higher than what it costs for INTERVENE to stop the massacres. So, even though the people of ATROCITY would prefer that the government of INTERVENE pursue both the just and

unjust aims rather than pursue neither, and even though they would contract accordingly, it does not necessarily follow that the INTERVENE is guilty of no wrong for taking the borderland.

Despite the fact that, for INTERVENE, the cost of humanitarian intervention is significantly lower than the value of the borderland, such intervention has a cost nonetheless. As a result, it might be the case that INTERVENE is not obligated to provide humanitarian intervention. This, of course, depends on our views regarding supererogation and positive rights. I will not explore these issues here. Instead, I will leave open the possibility that it is morally permissible for INTERVENE not to assist ATROCITY.

The conclusion that INTERVENE's war is narrowly unjust can be generalized. An element of a war is unjust only if those waging the war could have refrained from pursuing the unjust element. Any war with an unjust element that is not necessary for the achievement of just aims is itself unjust, since those who wage the war could have refrained from pursuing the unjust element. This means that a war of self-defense against a hostile, implacable aggressor unresponsive to diplomacy would be unjust if the defending nation included a small unjust element in their war. It might seem absurd to categorize both the aggressor's war and the defender's war as unjust -- this categorization elides significantly moral differences in the respective wars. However, I have presented a more fine-grained categorization of unjust wars that distinguishes wars that are mostly good from wars that are mostly bad -- the defender's war is narrowly unjust, where as the aggressor's war is broadly unjust.

With this conclusion in hand, I can now show how the account of individual responsibility for collective action affects the responsibility of combatants who participate in the most common type of war -- one with just and unjust aims.

8.5. RESPONSIBILITY IN MORALLY HETEROGENEOUS UNJUST WARS

Returning to the previous example, call the war that INTERVENE fights against ATROCITY 'W'. This war consists of two aims -- Φ and Ψ . The former is the aim of stopping the massacres. The latter is the aim of annexing the borderland. Each of these aims at least partly consists of operational objectives, which in turn, partly consist of tactical objectives, which partly consist of missions. Suppose that the unjust aim of annexing the borderland is partly composed of a just operational objective, e.g., rescuing refugees displaced by the pursuit of Φ . This just operational objective, Ψ_A (necessarily) consists of just tactical objectives and missions. Also suppose that the unjust operational object partly composing the unjust aim of the war includes a few just tactical objectives and missions. The following diagram depicts the relations among the elements of war W, as well as the moral status of each element:

war	aim	operational objective	tactical objective	mission
W unjust	Φ just	ΦA just	ΦA1	ΦA1 _a just
			just	ΦA1 _b just
		ΦB just	ΦA2	ΦA2 _a just
			just	
		ΨA just	ΦB1	ΦB1 _a just
			just	ΦB1 _b just
	Ψ unjust	ΨA just	ΨA1	ΨA1 _a just
			just	ΨA1 _b just
			ΨA2	ΨA1 _c just
			just	ΨA2 _a just
		ΨC unjust	ΨC1	ΨC1 _a just
			just	
			ΨC2	ΨC2 _a just
			unjust	ΨC2 _b unjust
	ΨC3	ΨC2 _c unjust		
	unjust	ΨC3 _a unjust		
		ΨC3 _b unjust		

Though an unjust war can consist partly of just elements, no just element can consist of unjust subsidiary elements. This is for the same reason that no just war can consist of unjust elements: any such element would constitute a harm unnecessary for accomplishing the just elements of the war. So although the unjust aim Ψ consists of just and unjust subsidiary elements, the just aim Φ consists solely of just subsidiary elements. Likewise, the just tactical objective, $\Psi C1$, though it is subsumed by an unjust operational objective, can itself subsume only just missions.

In what follows, I consider individual responsibility for participating in the just elements of an unjust war, first according to a non-collectivist account of individual responsibility, then according to the account I have developed so far.

8.5.1. AN INDIVIDUALISTIC ACCOUNT OF INDIVIDUAL RESPONSIBILITY

Consider a combatant, *c*, who participates in mission $\Psi C2a$. This morally just mission is partly constitutive of an unjust tactical objective, which is partly constitutive of an unjust operational objective, which is, in turn, partly constitutive of an unjust aim. If we ignore secondary responsibility, participating in $\Psi C2a$ is just. In such a case, the combatant, *c*, is contributing to a just mission. This just mission, $\Psi C2a$, partly composes an unjust tactical objective, $\Psi C2$; if it is not *because* of $\Psi C2a$ that $\Psi C2$ is unjust, then though *c*'s contribution to $\Psi C2a$ entails a contribution to $\Psi C2$, *c* does not contribute to what makes $\Psi C2$ unjust. Put differently, though *c* contributes indirectly to $\Psi C2$, *c* does not contribute to the 'unjust parts' of $\Psi C2$. It is difficult to see, then, how *c*'s participation in $\Psi C2a$ could be unjust, assuming we ignore secondary responsibility.

But what if the occurrence of $\Psi C2b$ and $\Psi C2c$ -- both of which are unjust -- depend on $\Psi C2a$? If $\Psi C2$ is unjust because of $\Psi C2b$ and $\Psi C2c$, and these two missions depend on $\Psi C2a$, then it seems that it is at least partly because of $\Psi C2a$ that $\Psi C2$ is unjust. But in such a case, $\Psi C2a$ would not be just -- *contra* what is assumed *ex hypothesi*. Put more generally, if an element of a war is unjust, and that element has a just sub-element, then

this just sub-element is not a necessary or sufficient condition for the occurrence of the unjust element.

The upshot is that, on a non-collective account of individual responsibility, a combatant is not responsible for the unjust elements of a war, if the combatant contributes solely to the just elements of the war. But we can go further and say that on such an account, it is often permissible for a combatant to participate in the *unjust* elements of a war. I claimed in section 8.2. that whether a set of aims (and thus a war) is just depends on what alternatives are available to the government. In the example I have given, the government of INTERVENE has the option of waging a war with only just aims. But unlike a government, individual combatants typically do not have the power to choose what aims a war will have. Because an agent can be morally required to do only what that agent is capable of doing, a typical civilian cannot be morally required to change the aims of an unjust war fought by her government. A combatant, in the best of circumstances, can either choose to participate or not to participate.

Does the fact that to promote a narrowly unjust war is to promote an unjust war provide a decisive reason for the combatant not to promote the narrowly unjust war? It is hard to see how the fact that the war as a whole is unjust precludes the permissibility of participation in the just aim in that war. The fact that the war is unjust is compatible with the claim that killing the combatants who are participating in the massacres satisfies the constraints of discrimination, proportionality (narrow and wide), and necessity. But I go further and claim that it is also often permissible for combatants to participate in the *unjust*

aims of the war, since the unjust aims are, *for the combatants*, instrumental to the achievement of the just aims.

Put generally, whether a particular act is necessary for the achievement of a desired end may depend on whether we adopt a first-person or third-person perspective with respect to that act. For the government of INTERVENE, annexing the borderland is not subsidiary to the aim of stopping the massacres. Rather, the government has *chosen* to pursue the former if and only if it pursues the latter. But things are different for those combatants uninvolved in the government's choice of aims. The aims of the war are a matter of choice for the government but are facts about the world from the point of view of the combatant. Even though the aim of annexing the borderland is not subsidiary to the aim of stopping the massacres, the annexation of the borderland is -- for the combatant -- required in order to stop the massacres. The just and unjust aims are, for the combatant, packaged together. For the government of INTERVENE, annexing the borderland is not subsidiary to the aim of stopping the massacres -- not so, for the combatants fighting on the side of INTERVENE.

But does the fact that averting the massacres requires annexing the borderland -- which is a rights-violation -- provide decisive reason not to participate in the annexation?

Sometimes it is permissible to infringe rights if doing so is necessary to avert significantly worse consequences, such as massacres, which is also a far more egregious type of rights-violation than the infringement of territorial sovereignty. It is impermissible for the government of INTERVENE to annex the borderland because doing

so is *not* necessary to avert the massacres. But for combatant choosing between the narrowly unjust war and no war at all, those rights infringements *are* necessary to stop the massacres. Hence it is permissible to participate in the unjust aims of a narrowly unjust war when the only other option is to promote no war at all. Put differently, if W in figure 1 is a *narrowly* unjust war, then it is possible that it is permissible to participate in Ψ , or any of the unjust elements subsidiary to Ψ -- assuming a non-collectivist account of individual responsibility.

8.5.2. THE COLLECTIVIST ACCOUNT OF INDIVIDUAL RESPONSIBILITY

If we include secondary responsibility, does participating in $\Psi C2a$ remain permissible? After all, those participating in $\Psi C2a$ are also participating in $\Psi C2$, which is unjust, as is ΨC , Ψ , and W . If a combatant bears secondary responsibility for the contributions made by those not only in $\Psi C2a$, but in the rest of W as well, then the combatant bears more responsibility for unjust than just contributions to the unjust war. As a result, though $\Psi C2a$ is a just mission, participating in it is impermissible. But if participating in $\Psi C2a$ is impermissible, in what sense is it a just mission? It is a just mission in that the mission, isolated mereologically from the war of which it partly composes, promotes the well-being of others without contributing to any harms or injuries. But a combatant, c , who participates in such a mission, also participates in an unjust war, and thus bears secondary responsibility for the acts committed by combatants who contribute to unjust missions, even if the combatant does not participate in any unjust missions.

So once we consider the affects of secondary responsibility, it seems unlikely that *c* is morally permitted to participate in even the just aims of an unjust war.

But this conclusion is premature, for several reasons. Though everyone in the collective bears secondary responsibility for the contributions made by everyone else in the collective, a combatant bears more secondary responsibility for the contributions made by the sub-collectives of which the combatant is a part. Recall from section 7.5.1 that if combatants share contributory intentions, then the degree of secondary responsibility that each bears for the other's action is augmented; and if they also share a mutual desire for the achievement of their respective functional role's shared object, then secondary responsibility is further augmented. Suppose that *c* shares contributory intentions with other members of $\Psi C2a$ to contribute to this just mission. And *c* desires the achievement of this mission. Moreover, *c* does not have a particular intention to contribute to $\Psi C2$, ΨC , or Ψ -- nor does he desire the achievement of any of these aims. As a result, *c* will bear greater secondary responsibility for $\Psi C2a$ than he does for the other elements of the war. So if *c* participates in $\Psi C1a$ then he is more likely to bear greater just over unjust responsibility for the war than would be the case if she participated in $\Psi C2a$. And if *c* participates in $\Psi C2a$, she is more likely to bear greater just over unjust responsibility for the war than would be the case if she participated in $\Psi C2c$ or $\Psi C3a$, for example.

There is another reason why it might be permissible to participate $\Psi C1a$, even given the effects of secondary responsibility. Recall from section 8.4 that $\Psi C2a$ cannot be necessary for the achievement of $\Psi C2$; otherwise $\Psi C2a$ could not be a just mission.

Moreover, $\Psi C2a$ cannot causally contribute to $\Psi C2$ without jeopardizing the just status of the former -- otherwise, a contribution to $\Psi C2a$ would result in a contribution to $\Psi C2$. $\Psi C2a$ is at best mereologically rather than causally related to $\Psi C2$, in that the former is a proper part rather than a cause of the latter. If $\Psi C2a$ does significant good, then the primary responsibility that c bears for the contributions to $\Psi C2a$ might outweigh the secondary responsibility that he bears for the unjust elements of the war.

Note that if a just mission causally contributes to a just tactical objective, as is the case in $\Psi C1a$, then c bears some primary responsibility not only for the just mission, but for the just tactical objective as well. Note that for c to bear any *primary* responsibility for the tactical objective (or other elements of the war, including the war itself), c must be causally rather than merely constitutively related to these elements of the war. Whether the just mission causally contributes to a just tactical objective (or any other element of the war), is itself irrelevant to the degree of *secondary* responsibility that c bears for that element. After all, c is a fellow member of those participating in the just tactical objective whether the just mission in which c participates is causally or constitutively related to the unjust tactical objective.

So far I have mentioned two ways that responsibility for $\Psi C2a$ might outweigh secondary responsibility for the unjust elements of a war. The first involved augmenting secondary responsibility for $\Psi C2a$ relative to the other elements of the war. The second involved increasing the good effects of $\Psi C2a$, in which case primary responsibility for $\Psi C2a$ weighs more heavily in favor of participating in the mission. But both involve comparing

the magnitude of a combatant's responsibility for a just element of a war with the magnitude of a combatant's responsibility for an unjust element of a war. But at one remove this tactic is hopeless as a means of demonstrating that even given secondary responsibility it is possible to permissibly participate in a just element of an unjust war, since the fact that an act involves net positive responsibility does not entail the permissibility of that act.

More specifically, it is impossible for *c* to bear secondary culpability for Ψ , *and* permissibly contribute to $\Psi C1a$ or $\Psi C1$ (or to any other element in *W*) qua member of *W*. This is in spite of the fact that, *ex hypothesi*, $\Psi C1a$ and $\Psi C1$ are permissible to pursue in isolation from the rest of the war. By *reductio*, suppose that *c* bears *secondary* culpability for Ψ *and* permissibly contributes to $\Psi C1a$ or $\Psi C1$. Contributing to $\Psi C1a$ and $\Psi C1$ qua member of *W* presupposes membership in *W*, which entails secondary responsibility for Ψ . A permissible contribution to $\Psi C1a$ and $\Psi C1$, made qua member of *W*, cannot entail blame for a wrongdoing (bad motives aside). If the permissible act involves saving the lives of many at the cost of injuring a few, then perhaps the actor might be *liable* to compensate those he permissibly injured. But he is not morally blameworthy for the injury, assuming of course the injury was unavoidable and sufficiently minor. The basis of his alleged culpability for Ψ (among other events) is his voluntary act of joining *W*. It seems that either *c* is not at all to blame for the unjust acts committed by the collective of which he is a part, or *c* cannot permissibly contribute to Ψ . This exclusive disjunction holds even if *c*'s primary and secondary responsibility for $\Psi C1a$ and $\Psi C1$ outweighs his secondary responsibility for the unjust acts that *W* commits. If the responsibility here is

culpability, and *c* bears *any* culpability for the unjust acts committed by *W*, then *c* cannot permissibly contribute $\Psi C1a$ and $\Psi C1$.

In summary, though secondary responsibility widens the radius of a combatant's responsibility, it is still nonetheless permissible, sometimes, to participate in a just aim of an unjust war, provided that the combatant intends to contribute only to the just aim. If we take solely primary responsibility into account, it will often be permissible to participate in the unjust aims of a narrowly unjust war. But if we also take secondary responsibility into account, it is never permissible to participate in an unjust aim of an unjust war, regardless of whether the war is narrowly or broadly unjust.

Bibliography

- Arendt, H. 1987, "Collective Responsibility," in James Bernauer (ed.), *Amor Mundi*, Dordrecht: M. Nijhoff.
- Bates, S. 1971, "The Responsibility of 'Random Collections'," *Ethics*, 81: 343.
- Bratman, M. 1987, *Intentions, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
 ----. 1992, "Shared Cooperative Activity," *Philosophical Review*, 101/1: 327.
 ----. 1993, "Shared Intention," *Ethics*, 104: 97-103.
- Cooper, D. 1968, "Collective Responsibility," *Philosophy*, 43: 258-268.
 ----. 1998, "Responsibility and the 'System'," in Peter French (ed.), *Individual and Collective Responsibility*, Rochester, VT: Schenkman.
- Copp, D. 1979, "Collective Actions and Secondary Actions," *American Philosophical Quarterly*, 16: 177-186.
 ----. 1984, "What Collectives Are: Agency, Individualism and Legal Theory", *Dialogue*, 23: 253-268.
- Corlett, J. A. 2001, "Collective Moral Responsibility," *Journal of Social Philosophy*, 32: 573-584.
- Davidson, D. 1963, "Actions, Reasons, and Causes," reprinted in *Essays on Actions and Events*, Oxford: Oxford University Press, 1980, pp. 3–20.
 ----. 1978, "Intending," reprinted in *Essays on Actions and Events*, Oxford: Oxford University Press, 1980, pp. 83–102.
- Downie, R. S., 1969, "Collective Responsibility", *Philosophy* 44: 66.
- Feinberg, J. 1968, "Collective Responsibility," *Journal of Philosophy*, 65: 674.
 ----. 1970, "Collective Responsibility," in *Doing and Deserving: Essays In The Theory of Responsibility* Princeton: Princeton University Press, pp. 222- 251.
- French, P. A. 1979, "The Corporation as a Moral Person," *American Philosophical Quarterly*, vol. 16, pp. 207-217.
 ----. 1984, *Collective and Corporate Responsibility*, New York: Columbia University Press.
 ----. 1995, *Corporate Ethics* Fort Worth: Harcourt Brace College Publishers.
 ----. ed. 1998, *Individual and Collective Responsibility*, Rochester, VT: Schenkman.
- Gilbert, M. 1989, *On Social Facts*, New York: Routledge.
 ----. 1994, "Remarks on collective belief," in Frederick Schmitt ed. *Socializing Epistemology*. Maryland: Rowman & Littlefield.
 ----. 1997, "Group Wrongs and Guilt Feelings," *Journal of Ethics*, 1: 65-84.
 ----. 2000, *Sociality and Responsibility*, Lanham, MD: Rowman and Littlefield.
- Graham, K. 2000, "Collective Responsibility," in *Moral Responsibility and Ontology*, Ton Vanden Beld (ed.), Dordrecht: Kluwer.
- Held, V. 1970, "Can a Random Collection of Individuals be Responsible?", *Journal of Philosophy*, 67: 471-481.
- Hurka, T. 2002, "Proportionality in the Morality of War," *Philosophy and Public Affairs* 33 at pp. 42.
- Jaspers, K. 1961, *The Question of German Guilt*, trans. by E.B. Ashton, New York: Capricorn.

- Keeley, M. 1981, "Organizations as Non-Persons," *Journal of Value Inquiry*, vol. 15, pp. 149-155.
- Kutz, C. 2007, *Complicity: Ethics and Law for a Collective Age*, Cambridge University Press.
- Ladd, J. 1984, "Corporate Mythology and individual Responsibility," *International Journal of Applied Philosophy*, vol. 2, no. 1.
- Lewis, D., 1989, "The Punishment that Leaves Something to Chance," *Philosophy and Public Affairs*, vol. 18, pp. 53-67.
- May, L. and Hoffman, S. (eds.), 1991, *Collective Responsibility: Five Decades of Debate in Theoretical and Applied Ethics*, Savage, MD: Rowman and Littlefield.
- May, L. 1983, "Vicarious Liability and Corporate Responsibility," *Philosophical Studies*, vol. 43 pp. 69-82.
 ----. 1987, *The Morality of Groups*, Notre Dame: University of Notre Dame Press.
 ----. 1992, *Sharing Responsibility*, Chicago: University of Chicago Press.
- McGray, H. 1986, "Morality and Collective Liability," *Journal of Value Inquiry*, 20: 157-165.
- McMahan, J. 2004, "The Ethics of Killing in War," *Ethics* 114: 693-733.
- McMahan, J. 2005, "The Basis of Moral Liability to Defensive Killing," *Nous Supplement* 15 (2005): 386-405.
- McMahan, J. 2007, "Precis: the Morality and Law of War," *Israel Law Review* 40: 310-323.
- McMahan, J. 2007, "On the Moral Equality of Combatants," *Journal of Political Philosophy* 14: 377-93.
- Meijers, A. 1994, *Speech Acts, Communication, and Collective Intentionality: Beyond Searle's Individualism*, Leiden.
 ----. 1999, *Belief, Cognition, and the Will*, Tilburg: Tilburg University Press, 59-71.
 ----. 2003, "Can Collective Intentionality be Individualized?," *American Journal of Economics and Sociology* 62, 167-93.
- Narveson, J. 2002, "Collective Responsibility," *Journal of Ethics*, 6: 179-198.
- Parfit, D. 1986, *Reasons and Persons*, Oxford University Press, USA, pp. 67-86
- Pettit, P. 2003, "Groups with Minds of their Own," in Schmitt F. (ed) *Socializing Metaphysics*, Rowman and Littlefield, pp. 167-93.
- Radzik, L. 2001, "Collective Responsibility and Duties to Respond," *Social Theory and Practice*, 27: 455-471.
- Raikka, J. 1997, "On Dissociating Oneself from Collective Responsibility," *Social Theory and Practice* 23: 1-9.
- Risser, D. T. 1978, "Power and Collective Responsibility," *Kinesis*, vol. 9, no. 1 (1978) pp. 23-33.
 ----. 1996, "The Social Dimension of Moral Responsibility: Taking Organizations Seriously," *Journal of Social Philosophy*, vol. 27, no. 1, pp. 189-207.
- Searle, J. 1990, "Collective Intentions and Actions," In *Intentions in Communication*, P.Cohen, J. Morgan, and M.E. Pollack, eds. Cambridge, MA: Bradford Books, MIT press.
 ----. 1995. *The Construction of Social Reality*. New York, N.Y.: Free Press.

- Scanlon, T.M., 1990, "Promises and Practices", *Philosophy and Public Affairs*, 19(3): 199–226.
- Thompson, J. J. 1987, "The Decline of Cause," *The Georgetown Law Journal*, vol. 76, pp. 137-150.
- Tollefsen, D. 2002, "Organizations as True Believers," *Journal of Social Philosophy*, vol 33 (3): pp. 395-411.
- Tuomela, R. 1991, "We Will Do It Again: An Analysis of Group Intentions," *Philosophy and Phenomenological Research*, 60: 249-277.
- 1992, "Group Beliefs," *Synthese*, 91: 285-318.
- 2004, "*We-intention Revisited*," *Philosophical Studies* (125:3), 2005, pp. 327-369.
- Valasquez, M. 1983, "Why Corporations Are Not Morally Responsible for Anything They Do," *Business and Professional Ethics Journal*, vol. 2, no. 3.
- Velleman, D. 1997, "How to Share an Intention," *Philosophy and Phenomenological Research* LVII: 29-50.
- Wall, E. 2000, "The Problem of Group Agency," *The Philosophical Forum*, vol. 31, no. 2, pp. 187-196.
- Williams, B. 1982, "Moral Luck" in *Moral Luck*, ed., Bernard Williams, Cambridge, Cambridge University Press, pp. 20-39.
- Zimmerman, M. J. 1985, "Sharing Responsibility," *American Philosophical Quarterly*, vol. 22, pp. 115-122.

Curriculum Vita - Saba Bazargan

EDUCATION

Doctorate of Philosophy
Spring 2010
Rutgers University

Bachelor of Arts, Departmental Honors
Spring 2001
University of California at Los Angeles

OCCUPATION

Assistant Professor of Philosophy

PUBLICATIONS

Review of *Just and Unjust Warriors: The Moral and Legal Status of Soldiers*.
Eds. David Rodin and Henry Shue, *Ethics* 119 (April 2009)

“Social Contract.” and “Subjective Aim.” *American Philosophy: An Encyclopedia*. Eds. John Lachs; Robert Talisse. New York: Routledge, 2007