This is the final draft of Chapter One of
Group Duties: Their Existence and Their Implications for Individuals (OUP, 2019).
For citation purposes, the definitive version can be purchased at
https://global.oup.com/academic/product/group-duties-9780198840275?cc=at&lang=en&

Chapter One: The Tripartite Model

1.1 Introducing the Tripartite Model

We treat groups as if they have duties. To find evidence of this, you just need to open a newspaper. In January 2019, an editorial in The Observer addressed its readers in Britain: 'As a nation, we may have different views on immigration, but surely we can all agree that Britain has a moral duty to provide safe harbour to more people fleeing conflicts' (Observer Editor 2019). In March 2017, conservative commentator James Bickerton argued that 'conservatives have a particular duty to defend liberal-democratic civilization from Trump's authoritarianism' (Bickerton 2017). In October 2012, the former United Nations Secretary-General Ban Ki-moon told an audience at Drake University that '[t]he international community has a moral responsibility, a political duty and a humanitarian obligation to stop the bloodbath and find peace for the people of Syria' (UNDPI 2012).

Such statements are not readily understandable as asserting that each member of the relevant group has the duty. Each member of Britain cannot ensure that Britain provides safe harbour. Not even the Prime Minister can do this, if acting alone. So the duty to provide safe harbour is not a duty that any individual member of Britain can have. Likewise, no conservative can defend the entirety of liberal-democratic civilization from anyone's purported authoritarianism. Defending a civilization is simply not something any individual can do. Civilizations can be defended and upheld only by the combined actions of many individuals. So Bickerton's claimed duty must be held by conservatives as a group. And each member of the international community cannot, on its own, find peace for the people of Syria. This is true regardless of whether we take the members of that community to be individuals or states. No singular individual or state can act unilaterally to find peace for the people of Syria. It would take collaboration between many such entities to find peace. It would seem, then, that all three duties are being attributed to groups 'in their own right': the duty is being attributed to a group as a singular entity—or perhaps to the members as a plurality—not to each of the group's members taken severally.

This book asks: Do such attributions make conceptual sense? And what does that imply for the groups' members? In particular: when such attributions do make sense and are true, what do these group-held duties imply about the duties of the group's members? And when such attributions do not make sense, how should we charitably interpret such talk? What duties might such talk be taken to impute to the group's members? To answer these questions, we need a model that can take in details about the real-world groups targeted by such talk, and produce conclusions about (i) which of these real-world groups, if any, can have duties and (ii) what these groups' duties—or their inability to bear duties—imply about the duties of the groups' members.

This book develops and defends such a model. I call it the 'Tripartite Model' of group duties because it divides groups into three fundamental categories. Each category receives different answers to the above questions. To get a flavour of the three categories, reconsider the above examples.

The duty to ensure adequate air quality was attributed to Britain. In context, 'Britain' clearly refers to the UK (rather than Great Britain, which excludes Northern Ireland). The UK is a highly organized group. It has a decision-making procedure, which it operates in a more-or-less rational way. This procedure encompasses (amongst other things) the houses of parliament, ministerial departments, and judicial system. The UK's decision-making procedure subsumes the more specific procedures of these entities, which are used, respectively, to create, enforce, and apply the UK's laws. Via these procedures, the UK has the capacity to attend to moral considerations—even if it often fails to do so.

The outputs of the UK's overarching decision-making procedure—and the more specific procedures that it subsumes—distribute roles amongst members, where the roles are sufficient for enacting the UK's decisions. Roles include things like 'payer of the top tax rate', 'motorist who drives on the left', 'teacher of this particular curriculum'. These are all roles that are distributed by the UK's decision procedure, via its various sub-procedures. When such roles are enacted, their bearers are enacting the UK's decisions. What's more, each member of the UK is permitted, by the UK's decision-making procedures, to have inputs into the UK's decision-making procedure—whether those inputs take the form of votes, petitions, statements in parliament, evidence given to judges, and so on. The UK is (what I'll call) a collective.

In the second quote above, James Bickerton attributed a duty to conservatives. Importantly, we know he wasn't addressing a specific political party, since he was addressing a group that spans at least two countries (the UK and the US). So he wasn't addressing the UK's Conservative and Unionist Party—which, like the UK, is a collective. Instead, he was addressing the group that includes all people who hold certain values, goals, and beliefs. In particular, Bickerton implies that each conservative holds 'liberal-democratic-capitalist values', the goals of 'tax cuts and reduced business regulation', and the belief that 'a demagogue with a clear disregard for democratic norms . . . is not an acceptable candidate' (Bickerton 2017).

As explained above, Bickerton is not charitably understood as saying that each of these individual conservatives must defend liberal-democratic civilization. Instead, he was addressing the group comprised of all those who share these values, goals, and beliefs, and who are thereby disposed to interact with one another in such a way that those goals get achieved, and those values get respected, in ways that are constrained and guided by the beliefs. This group lacks the kind of group-level decision-making procedure, and the distribution of roles for enacting group decisions, that we find in collectives such as the UK. But its members share goals and are disposed to act positively towards one another in pursuit of those goals. Conservatives are a coalition.

Finally, Ban Ki-moon attributed an obligation to the international community. There is dispute about what, exactly, 'the international community' refers to, with some commentators suggesting that it refers to 'the west, of course, nothing more, nothing less' (Jacques 2006). But in the mouth

of Ban Ki-moon, speaking in his role as UN Secretary-General, it is more charitable to view 'the international community' as referring to that group of political states that is made up of members of the UN. To be clear, then, the international community is a group whose members are also groups. (To be specific, its members—states—are collectives.) But it's not clear that Ban meant 'the UN' when he said 'the international community'[1]—after all, the UN (whether General Assembly, Security Council, or another branch) does not have the authority to require its members to 'bring peace to Syria'. Let's assume that if he'd meant 'UN', he would have said 'UN'.

Instead, Ban can be interpreted as attributing a duty to the group that is composed of all those entities (states) that happen also to be members of the UN. He was not addressing this group qua UN. Instead, he was addressing this group simply as a collection of states, each of which has different values, goals, and beliefs, and each of whom can act outside the UN (though, he might have hoped, with the UN's blessing) in ways related to peace in Syria. The international community is neither a collective nor a coalition. Because the international community is neither a collective nor a coalition, it's a combination.

Collectives, coalitions, and combinations are the organizing categories of the Tripartite Model. In the most general and abstract terms: collectives are constituted by agents that are united under a rationally operated group-level decision-making procedure that has the potential to attend to moral considerations; coalitions are constituted by agents who each hold a particular goal and are disposed to work with the others to realize the goal, while lacking a group-level decision-making procedure that has the potential to attend to moral considerations; and combinations are constituted by any collection of agents that do not together constitute either a collective or a coalition. These three categories of group are exhaustive and mutually exclusive.

The Tripartite Model gives the following verdicts. First, neither coalitions nor combinations can have duties in their own right, that is, duties that are not identical to a conjunction of members' duties, or (equivalently) duties that are not mere shorthand for a conjunction of members' duties. I argue for this in Chapters 2 and 3. Chapter 2 challenges six arguments that purport to show that combinations and coalitions can bear duties. Chapter 3 tackles the positive task of providing arguments in favour of the idea that they cannot bear duties.

How, then, should we charitably interpret commentators such as Bickerton and Ban, when they attribute duties to coalitions and combinations? Chapter 4 argues that they should be understood as saying that each member of the relevant group has a duty to act in response to the others, with a view to defending liberal-democratic civilization or bringing peace to Syria. Chapter 4 gives these 'coordination duties' a precise structure and a set of sufficient conditions on their existence.

Coordination duties can be held by members of both coalitions and combinations. Yet Chapter 5 argues that coordination duties operate differently in coalitions, as compared with combinations. Specifically: when one has a coordination duty as a member of a coalition, that duty demands that one engage in a particular variant of 'we-reasoning', which I call 'coalition-reasoning'. By contrast, when held as a member of a combination, coordination duties demand that one engage in

1 I thank Toni Erskine for pressing this interpretation.

'I-reasoning'. (This distinction draws on Bacharach (2006), Tuomela (2007; 2013), Sugden (2003), Gold and Sugden (2007), and others.) Thus, the implications of (the lack of) group duties for members are different in coalitions and combinations. This affects not just how the members should reason, but also how they should act. As we'll see, it means coordination duties have more chance of producing valuable outcomes when they are held by members of a coalition than when they are held by members of a combination (all else being equal).

Chapter 6 defends my account of collectives, which is more permissive than other prominent accounts in the literature. That is, it's easy for a group to be a collective. Specifically, unlike Peter French (1984), I do not require collectives to have unchanging corporate policies; unlike Carol Rovane (1998), I do not require collectives to have a unified project; and unlike Christian List and Philip Pettit (2011), I permit dictatorial collectives, unanimity-based collectives, and other collectives in which the group's attitudes on a particular issue are determined by certain members' attitudes on that issue (so long as the collective can remain rational).

Chapter 6 argues that collectives can have duties in their own right. So, when a collective bears a duty, its members' duties are not best understood as coordination duties. Instead, when a collective bears a duty, this implies that each of its members has (what I call) a 'membership duty'. These are duties to act within and because of one's role in a collective, with a view to the collective doing its duty. In Chapter 7, I explain why membership duties follow from collectives' duties and how membership duties are structured. As will become clear, they're quite different from coordination duties.[2]

The Tripartite Model is summarized in the following diagram. The diagram also indicates which chapter argues for which aspect of the model.

2 A note on terminology. I will sometimes talk about 'members' and sometimes about 'individuals'. In both cases, I refer to agents. Members are agents who constitute a group; individuals are agents who may or may not constitute a group. Chapter 6 will argue that collectives are agents. So, sometimes, an agential constituent—member—of a group will be a collective. For example, the House of Commons is a collective agent that is a member of the collective agent that is the UK. The House of Commons is treated as an individual, when it comes to considering the implications the UK's duties has for its members. So, in general, 'individual' means 'individual agent', not 'individual human'.

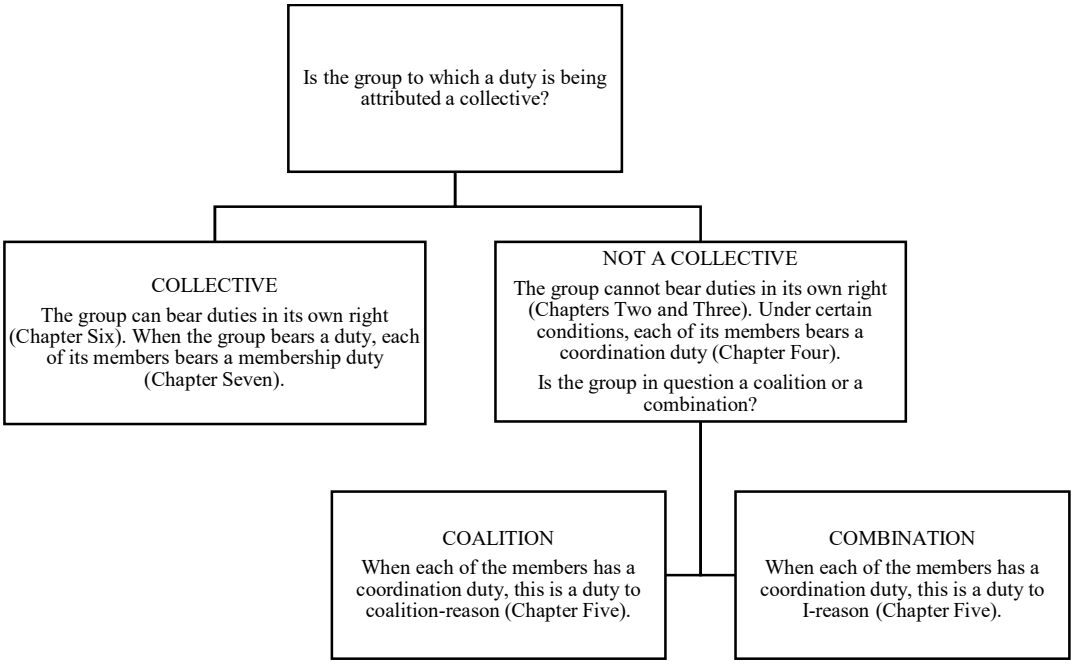

The rest of this introductory chapter will: first, relate group duties to some nearby philosophical issues; second, elaborate on the three types of group in more detail; and third, explain the scope of the book's aims.

1.2 Nearby Issues

Philosophers have addressed issues closely related to group duties. But we lack a comprehensive theoretical framework for group duties—where, by 'comprehensive theoretical framework', I mean a framework that can take in the facts about any group and produce answers to the questions of (i) whether that group can have duties in its own right and (ii) what that implies about the duties of the group's members.

To see that philosophers have addressed related issues, it helps to distinguish duty from responsibility. Those working on collective morality have addressed responsibility far more extensively than they have addressed duty. 'Responsibility' is a notoriously ambiguous and multifaceted concept (Hart 1968). We can distinguish at least three senses of group-level responsibility-for-an-outcome: causal, moral retrospective, and moral prospective (Collins forthcoming a). The last of these is what I mean by 'duty'.

Causal responsibility exists when one thing impacts or influences another. The causal responsibility of groups has been investigated extensively in the huge and long-running social scientific debate between methodological individualists and holists (cornerstones of which are Durkheim [1898] 1994; Weber 1922; Hayek 1942; Elster 1982; Kincaid 1986). Within that debate, the primary question is whether groups are eliminable from our explanations about what causes what within the social world. This debate obviously bears upon group duties. In particular, if groups' causal powers cannot be eliminated from our social ontology, then this gives us reason to

posit groups as distinct entities. We are then one step closer to attributing duties to groups, since groups' existence as distinct entities is plausibly necessary for them to have duties in their own right. But groups' distinct existence is not sufficient for their having duties in their own right; causal responsibility does not entail a capacity to bear duties. On the other hand, it might be that groups' duties must be posited within our moral theory, though their causal powers can be eliminated from our social scientific explanations. Groups would then be distinct entities, whose existence must be posited for moral purposes (for example, because their duties justify or ground individuals' duties)—but where groups are not posited because of their causal powers. So the causal-explanatory social-scientific question is not the one I'm concerned with.

There is another sense of group responsibility, which has received less attention than causal responsibility, but still more than duties. This is moral retrospective responsibility. This is the kind of responsibility at issue in debates about blame, free will, answerability, and accountability. It is sometimes called the 'basic desert' sense of responsibility (Pereboom 2001) and is often used interchangeably with 'blameworthiness' (Smiley 2010). I label this type of responsibility 'moral retrospective', to mark that it is essentially concerned with events that have already happened (unlike moral prospective responsibility, to be discussed shortly), and that it is concerned with moral evaluations of those happenings (unlike causal responsibility, which is a non-moral concept). If an entity has moral retrospective responsibility, then it is an appropriate target of reactive attitudes such as indignation, resentment, or gratitude. Numerous philosophers have developed theories of groups' blameworthiness or of individuals' blameworthiness in group contexts (Jaspers 1948; Feinberg 1968; Held 1970; French 1984; May 1987; May 1992; Kutz 2000; Gilbert 2002; D. Miller 2004; Thompson 2006; Tollefsen 2006; Pettit 2007; Mäkelä 2007; Smith 2009; Isaacs 2011; Strand 2012; Goodin and Lepora 2013).

Moral retrospective responsibility might seem tied to duties. It is tempting to think that blameworthiness and duty are two sides of the same coin: that an entity has a duty to produce some action or outcome if and only if that entity would be blameworthy if it did not produce that action or outcome (Lawford-Smith 2015). If so, then perhaps all and only the groups that are candidates for blameworthiness are candidates for duties; and perhaps groups' duties have implications for individuals that exactly parallel the implications that groups' blameworthiness has for individuals' blameworthiness.

But we shouldn't rush to assume that blameworthiness and duty are two sides of the same coin. Plausibly, they serve different functions in our lives. Chapter 3 will suggest that duty functions as a fitting input into an entity's first-person future-directed decision-making about what to do. By contrast, blameworthiness seems to function as a fitting output of second-personal past-directed judgements about what is appropriate (see, e.g., Smith 2015; Shoemaker 2015). If they serve such different functions, then it is an open question whether blame is fitting if and only if the target of the blame is capable of bearing duties. For this reason, we can't simply read our theory of group duties off our existing theories of group blameworthiness, even if we know one of those theories to be true.

What's more, even if blameworthiness and duty are two sides of the same coin, it's possible that the arguments in favour of some group-type being eligible (or ineligible) for duties will differ from

the arguments in favour of that group-type being eligible (or ineligible) for blameworthiness. What makes a creature duty-apt might be different from what makes it blame-apt—even if duty-aptness and blame-aptness rise and fall together. So we should explore groups' duties in their own right.

And even if groups' duty-aptness and blame-aptness have the same grounds, we should certainly explore group duties separately when it comes to implications for individuals: group duties may well 'distribute' differently than group blameworthiness. It seems prima facie plausible that a member's 'share' of a group's blameworthiness depends upon their personal contribution to the group's wrong—while a member's 'share' of a group's duty might depend, for example, on their power within the group. This can come apart from causal contribution to wrongs. To avoid making unwarranted assumptions on such issues, I will stay focused squarely on duties. I will refer to theories of group blameworthiness only when they have potential analogues regarding duties.

Other philosophers have come even closer to my target, since they have developed theories of group duties. But these theories have not been comprehensive. That is, they have not dealt with all types of group. We have theories about the duties of collectives and what collectives' duties imply for members, which I'll engage with in later chapters (e.g. Hindriks 2009; Copp 2007; Copp 2012). Most prominently in this category, we have the large business ethics corpus on corporate social responsibility—which deals with only one type of collective, namely, the for-profit limited-liability company, and which concerns itself with arguing about the substantive content of its duties (rather than the more general question of which kinds of company can have those duties and what those duties imply for members) (most famously, Friedman 1970). Also under the heading of collectives' duties, we have the much smaller debate on the moral duties of states and their implications for state members (Erskine 2001; Erskine 2010; Stilz 2011; Beerbohm 2012; Pasternak 2013; Lawford-Smith forthcoming), as well as the moral duties of intergovernmental organizations (Erskine (ed.) 2003). We also have theories of the duties of (what I call) combinations, in situations where the members of the combination happen to find themselves thrown into a morally significant situation alongside one another (Wringe 2014; Wringe 2016; Schwenkenbecher 2014; Björnsson 2014; Lawford-Smith 2015b). But none of these theories covers all three types of group.[3]

Finally, we have the well-established literature on shared action, joint agency, and group agency. This is highly pertinent for the Tripartite Model: in my terminology, philosophers have distinguished combinations from coalitions (e.g. Bratman 2014; Gilbert 1989; Sugden 2003; Tuomela 2007) and have distinguished combinations or coalitions from collectives (e.g. French 1984; Rovane 1998; Pettit and Schweikard 2006; List and Pettit 2011). But these theories do not attend to the distinctive issues involved in duty. (One almost-exception is Gilbert (1993; 2006), who connects plural subjecthood to the obligations of individuals, but without construing the latter as moral obligations.)

---

3 This isn't quite true, insofar as some philosophers might think we don't need a theory that differentiates random collections from highly organized groups. (Wringe and Björnsson, for example, might fall into this camp.) So these philosophers might take their theories of combinations' duties to be comprehensive across all groups. This book is devoted to arguing that such equal treatment of different group-types would be a mistake.

In sum, we don't yet have a theory that stitches these strands into one cohesive whole: a theory that provides a unified model for all three of combinations, coalitions, and collectives, where that model is focused specifically on these entities' capacity to bear duties. As we'll see, my arguments for the Tripartite Model will dispute others' arguments for many of the above-mentioned non-comprehensive theories.

1.3 The Tripartite Distinction

There are numerous ways of dividing up the logical space of groups (e.g. Isaacs 2011, 24–7; List 2014). Different divisions will be useful for different purposes. My purpose is to discover differences in groups' capacities to bear duties. So, my division between combinations, coalitions, and collectives is guided by when and why such differences arise. I don't claim that the tripartite distinction is the all-things-considered most important or fundamental way to categorize groups. But the coming chapters will argue that it is the most important and fundamental distinction for understanding groups' duties.

It's not always easy to categorize real-world groups under the three concepts. As with many concepts, there are likely to be vague or borderline cases. Not least, this is because many groups start off as combinations, then slowly become coalitions, then slowly become collectives. Consider, for example, protest movements. It's often impossible to pinpoint a single moment at which such an evolving group becomes a coalition, and then becomes a collective, and then perhaps falls back into being a combination as it achieves its goals. Still, the three types are not just three salient points on a continuum. This can be seen from the fact that there is no one metric with which we could label any such continuum. The three groups are different in kind—not merely in degree—even though there are examples on the cusp of each kind.

I'll start with the most complex type of group: collectives. This is because the other two group-types can be defined largely via the fact that they are not collectives. Although I take collectives first here, the book's substantive arguments will be clearest if the chapters work backwards through these types of group—starting with combinations and coalitions in Chapters 2 to 5, and dealing with collectives only in Chapters 6 and 7.

1.3.1 Collectives

A collective is constituted by agents that are united under a rationally operated group-level decision-making procedure that can attend to moral considerations. Some collectives are (arguably) also partly constituted by material infrastructure. But each collective is at least partly—and some are wholly—constituted by agents. The agents are united by a decision-making procedure. The procedure often includes many sub-procedures, including the informal, tacit, and vague procedures that we might subsume under the notion of a collective's 'culture'. In the most general terms, a 'decision-making procedure' takes in beliefs and preferences, and processes them to produce decisions. Rational operation requires that current decisions follow from current beliefs and preferences, and that current beliefs and preferences accord with past beliefs, preferences, and decisions plus any new evidence that has arisen since those past beliefs, preferences, and decisions were arrived at. In any decision-making entity—whether individual or collective—there is the possibility of conflict between (i) the entity's past beliefs and preferences, (ii) the entity's present

beliefs and preferences, and (iii) new evidence and inputs. If the procedure is to be rational, it must adjudicate these conflicts—deciding which beliefs and preferences will stay, which will go, and which will be deemed overridden in this instance. At the risk of stretching language, I'll assume that these adjudications are also 'decisions' that the procedure makes, so that the procedure can 'decide' on new beliefs and preferences, that then function as inputs into new decisions.

'Rational' operation does not require perfect rationality. This is clear in the case of individuals: I count as rational, yet it's far from true that all my beliefs, preferences, and decisions are fully coherent. But I aim to make my beliefs, preferences, and decisions consistent at one time, as well as consistent across times and circumstances. I mostly succeed. This level of coherence is, likewise, all that's required for a collective to count as rational.

What's more, I can attend to moral considerations. I often do so imperfectly, and I sometimes fail to do so at all. But I have a cognitive structure and material resources (e.g. a body, through which I can affect others) that enable me to consider and act upon moral considerations. Likewise, collectives have the basic structural and material resources to consider and act upon moral considerations. This is what I mean by 'can attend to moral considerations'. Chapter 6 will explain that this is not a high bar to meet, that the groups that meet it (i.e. collectives) are moral agents, and that they can therefore bear duties. The upshot is that almost all groups with decision-making procedures can bear duties: insufficient rationality, or incapacity to attend to moral considerations, are rare in groups with decision-making procedures.

Collectives meet three individually necessary and jointly sufficient conditions for the decision-making procedure to be group-level and for the agents to be united under that procedure. When these three conditions are met, the agents who are united under the procedure are the collective's members. So I'll refer to these agents as the 'members' in explaining the three conditions.

First, each member is committed (even if only tacitly) to abide by the procedure's results. This commitment can be overridden, but is presumptively decisive in the member's decision-making. To use the UK as an example: citizens' behaviour largely reflects a commitment to abide by the UK's decisions (that is, its laws). Of course, citizens sometimes break laws for reasons of personal conviction or individual gain. But as long as occasional law-breaking is the overriding of a general presumption of abidance on a member's part, they are committed. Now, we might think their commitment is made under duress, and therefore that their membership (insofar as it hinges on that commitment) should not generate any duties for them. But membership duties may well demand less of members, the more coerced their membership is.[4] Chapter 7 will have more to say about this.

Second, the beliefs and preferences that the procedure takes as inputs, and the way the procedure processes those inputs to form decisions, systematically derive from the behaviour (e.g. deference, votes, meeting contributions, etc) of members, while being operationally distinct from the inputs and processes that any member uses when deciding for herself. These aspects are operationally

4 Gilbert (2006) likewise allows for coerced membership duties, though (as noted above) she doesn't view these as moral duties.

distinct insofar as: first, the collective has the capacity to take in different considerations when forming its decisions than the members take in when forming their own decisions; second, its method for processing those considerations is different from the method of any one member when deciding for herself; and third, the decisions it produces are not the straightforward conjunction of members' decisions. (List and Pettit (2011) also emphasize this. Chapter 6 differentiates my account from theirs.)

To again use the UK as an example: the beliefs, preferences, and decisions of the UK systematically derive from the manifestos of political parties, the directives of ministers, the judgements of juries, and so on. The first two of these (at least) systematically derive from the votes, letters, petitions, and so on of members (which includes both citizens and residents). So the UK's decisions systematically derive from those of members.

At the same time, the UK's beliefs, desires, and decisions are operationally distinct: the UK takes in different inputs than any given citizen does when deciding for herself (in particular, because the UK aims to take in the inputs of all voters); the UK's methods for processing its beliefs and preferences to form decisions include parliamentary debates, Cabinet meetings, and so on—these are decidedly not the procedures that members use when processing their own beliefs and preferences to make their personal decisions; and, finally, the decisions of the UK are not just 'the Prime Minister decides this and the Foreign Secretary decides that and . . .' Instead, the UK resolves differences in members' decisions to arrive at an unequivocal collective position. (This is most clearly exemplified in the 'collective responsibility' of Cabinet (Brady 1999).)

The third jointly-sufficient condition on a decision-making procedure's being group-level, and of agents being united under the procedure, is that the enactment of the group's decisions requires actions on the part of the members, where those actions are also properly understood as attributable to the collective. (French (1984) also emphasizes this. Chapter 6 differentiates my account from his.) Suppose the UK decides that it will teach its schoolchildren a particular curriculum. The enactment of this decision requires actions on the part of thousands of members: teachers, regulators, curriculum designers, and so on. The actions of a given teacher, when teaching this curriculum, are their own. But they are also properly understood as attributable to the state. To see this, we can imagine that the state takes the position that evolutionary theory will be taught in biology classes. To a teacher who doesn't believe evolutionary theory, their action of teaching evolutionary theory is their own—but it is not properly attributable to just them. Instead, they are doing what the state demands; they are acting as a constituent, organ, part, or (as I will usually say) member of the state. Without referring to the state, we cannot locate the decision-maker that is (in large part) behind their action.

I'll say more about my view of collectives in Chapter 6. But the early chapters will be more intelligible if I illustrate my view's permissiveness now. This will help demonstrate which groups don't count as collectives and, so, which groups are the targets of Chapters 2 to 5.

To illustrate the permissiveness, imagine a group of three friends who are at the beach. Numerous decisions must be made: where to lay their towels, where to go for lunch, and so on. Such a group is probably composed of agents that are united under a rationally operated group-level decision-

making procedure that can attend to moral considerations. The procedure is probably conversation-based consensus. This procedure can become established simply by each member's taking a conversational and consensual stance to the various decisions—each asking the others where, for example, they would prefer to go for lunch, and why, until all agree. Such a procedure can be rationally operated, just so long as the group doesn't decide, for example, both to get burgers for lunch and not to get burgers for lunch. By taking part in the conversations (which can be done simply by staying silent to indicate indifference), each member tacitly commits to abide by the procedure's results.

The group has the capacity to have different beliefs and preferences from the members. For example, Laura might prefer sushi for lunch, but upon learning the preferences of the others, partakes in the consensus over burgers. We can even imagine a scenario where all members aim to please all the others, but all are wrong about what all the others want, so they end up using conversation-based consensus to go somewhere that none of them wants to go to. (This is perhaps more likely in England than elsewhere.) Here, the group prefers something (having pizza, say) that none of the members prefers. Likewise, such a group might believe something no member believes (as Gilbert (1987, 191) discusses). Such beliefs and preferences function as inputs into the group's decision—say, its decision to have pizza—even though those beliefs are different inputs from those of any member when deciding for herself.

Likewise, the group processes its beliefs and preferences differently than the members process their own beliefs and preferences: the group's method for making decisions is conversation-based consensus. Members do not use conversation-based consensus to form their own beliefs, preferences, and decisions. So the collective processes its inputs in a way none of the individuals does. And the group's decision is different from the conjunction of members' decisions. Even if all the members genuinely want pizza, it is only the group that decides 'the group will have pizza'. Each member merely decides 'I will endorse the group's having pizza'. So the group's decisions are not the conjunction of members' decisions.

Finally, the enactment of the pizza-for-lunch decision requires that members act (by walking to the pizza place). These actions are attributable to the group in the following sense: the explanation of why any one of the individuals walks to the pizza place requires referring to the fact that she's a member of a group that has decided to have pizza for lunch. The group has given her (and all the others) the role 'member who walks to the pizza place'. She walks there as part of her role in the group, because of her role in the group, because that role has been distributed to her by the group. She goes to the pizza place for lunch, but this action also partly constitutes the group's action of going there for lunch.

This permissive picture can be scaled up to larger-scale and more significant collectives. Thus, on my account, a collective might be a group whose members are states working together on a military invasion (in Collins (2014), I argued that 'coalitions of the willing' can be collectives), or whose members are corporations engaged in price-fixing. Again, Chapter 6 will say more about why we should hold this permissive account of collectives, and the sense in which their duties are irreducible to members' duties.

1.3.2 Coalitions

Members of coalitions meet three conditions: (1) each holds the same goal (preference, desire) as the others, where it's out in the open between them that the goal (preference, desire) is held by each (even though it might be held only tacitly or implicitly); (2) each is disposed to act responsively to the others (insofar as they encounter one another) to realize the goal (satisfy the preference/desire); but (3) the members are not united under a rationally operated group-level decision-making procedure that can attend to moral considerations. They might be united under an irrational or morally incapable procedure, but more likely they will be united under no group-level procedure at all.

When I say that it's 'out in the open' that the goal is endorsed by each, I don't mean that each member knows exactly who all the other members are. Coalitions include, for example, 'environmentalists', 'the oil lobby', 'democracy-promoting states', and—as we saw above—'conservatives'. Each environmentalist holds the goal of the environment being protected; each oil lobbyist holds the goal of oil-friendly public policy; and each democracy-promoting state holds the goal of democracy being promoted; and so on. This doesn't require that each environmentalist (lobbyist, state) knows who every other one is.

To be more precise: each member of the coalition might hold merely 'T-conditional common knowledge' that each other member holds the goal. Adapting Gold and Sugden's (2007, 131) definition of T-conditional common knowledge, I'll assume it's 'out in the open' that a goal G is endorsed by each member of some group T if: each member of T endorses G; each member of T knows there is more than one member of T; each member of T knows that, if some agent A is a member of T, then A endorses G; each member of T knows that, if some agent A is a member of T, then A knows that, if some agent A is a member of T, then A endorses G—and so on. In coalitions, it is also 'out in the open' that there is more than one member of T. When I say that each member 'knows' the facts at each level of iteration, I don't mean that each member has an explicit belief at each level. Rather, it'll usually be enough that there are no false beliefs at any level of iteration. (Here I follow Regan (1980, 129–30).)

The goal's being 'out in the open' in this way is crucial to making the difference between agents' duties in coalitions and their duties in combinations. To foreshadow: Chapter 5 will argue that assertions of coalitions' duties should be understood as assertions that members have duties to engage in a variant of we-reasoning that I call 'coalition-reasoning'. When an agent coalition-reasons, the agent presumes that other agents will do their parts in a certain pattern of actions, and then the agent does their own part in that pattern. For coalition-reasoning to be sensible, the agent must be able to 'hold fixed' that the other agents will do their parts in the collective pattern of actions about which the agent reasons. In Chapter 5, I will argue that this 'holding fixed' is justified in coalitions.

My definition of the 'out in the open' condition implies that, for example, Stephen Hawking and Isaac Newton are each part of a coalition, constituted by all agents who hold the goal 'scientific discovery' and are disposed to act responsively to one another (insofar as they encounter one another) to realize scientific discovery. This doesn't require that Hawking and Newton ever

encounter each other. This has the benefit of allowing that, say, conservatives in 1970 and conservatives in 2018 are part of the same coalition, which persists in time throughout changes in membership. As we'll see in Chapter 5, this will have upshots for how conservatives in 1970 should have conceived of their coordination duties: they had duties to engage in coalition-reasoning, not just with each other, but also with the conservatives of the future. This means they have a duty to consider a different (and larger) pattern of actions than if they did not we-reason with these future conservatives in mind. This, I suggest, is a virtue of the account.

A few more clarifications are necessary about coalitions. One is that, for the goal to be truly the same amongst the members, the goal must not be relativized to any of the members. For example, above, I mentioned 'democracy-promoting states' rather than 'democratic states'. This is because democratic states do not share any goal that is neutral between the states. Each democratic state has the goal of itself being democratic. For example, if New Zealand has the goal of 'democratic decision-making within New Zealand' and Australia has the goal of 'democratic decision-making within Australia', then they do not form a coalition (even if these goals are out in the open and so on). This is because these goals are agent-relative: the agent that holds the goal is a non-removable part of the goal itself. If two agents each hold an agent-relative goal with the same content (the content being, e.g., 'democratic decision-making'), then these are not the same goal.

This is important, because often agent-relative goals with the same content are mutually incompatible: I might have the goal 'Stephanie winning Stephanie and Holly's tennis match' and Holly might have the goal 'Holly winning Stephanie and Holly's tennis match'. These agent-relative goals are mutually incompatible, despite having the same content when framed as 'winning the tennis match'. So Holly and I are not in coalition around 'winning the tennis match'. By contrast, the goal of 'playing tennis' is neutral between myself and Holly, and when we play tennis we are in a coalition around this goal. Likewise, the goal of 'democracy being promoted' is agent-neutral: it does not irreducibly refer to the agent that holds the goal. Coalitions can be unified around only goals that each member of the coalition can, at the same time, endorse. Holly and I can be in a coalition around the goal 'Stephanie winning Stephanie and Holly's tennis match' (if Holly wants to lose), and Holly, Avia, and I can be allied around the goal 'Stephanie winning Stephanie and Holly's tennis match' (even though Avia is not playing in the match). But Holly and I cannot be allied around 'winning the tennis match' if we each wish this for ourselves.

This implies that some politically important groups—that we might intuitively think of as coalitions—are not really coalitions. Consider, for example, 'bankers' or 'the rich'. The members of these groups have politically significant features in common. But there is no common goal that they share (each banker might want only their own bank to be preserved, for example). Often, the members of these groups are in competition with each other. This will mean that their duties are not duties to coalition-reason. Again, this will be defended in Chapter 5.

Another clarification on condition (1) is that the goal might be ongoing and tacit. Consider families. Members of many families have the goal 'stick together' or 'be there for one another'. There is no precise time at which such goals are achieved; instead, the goal is an ongoing project in the minds of the members. Such a goal induces duties for members to engage in coalition-reasoning. The same is true of goals that are tacit. For example, there's a sense in which everyone

in this café shares the goal 'it not being the case that a bomb goes off in this café right now'. That goal is 'tacit' or 'implicit': no one in the café is thinking about the goal (except me). Yet if we were to acquire coordination duties qua members of this coalition—for example, if a bomb alarm started ringing—then we would each have a duty to coalition-reason.

What about condition (2)? When I say members of a coalition are 'disposed to act responsively to the others to realize the goal', I mean they are disposed to act upon each other with a view to realizing the goal. Chapter 4 will have more to say about responsiveness. We can analyse it further in various ways which place more or fewer demands on the members. For example, some coalitions are examples of Margaret Gilbert's (1989) 'plural subjects', in which the individuals have a 'joint commitment' to pursue their goal 'as a body'; or Michael Bratman's (1992, 328) 'shared cooperative activity', in which 'each participating agent [knows] that the other is attempting to be similarly responsive'; or Christopher Kutz's (2000) acting together through 'participatory intentions' to do one's part in an action or project together with others; or Raimo Tuomela's (2006, 38) joint intentional action, which 'amounts to the group members' jointly intending X and jointly acting on the basis of their joint intention'. Each of these specifications is more demanding than the minimal notion I will work with (defined by characteristics (1), (2), and (3) above). So, like collectives, coalitions are hugely varied in their complexity.

To illustrate how coalitions differ from collectives, consider again the group of friends at the beach. Near the group of friends, there might be various sunbathers who didn't come to the beach together. Once at the beach, these individual sunbathers might (though will not necessarily) form a coalition. For example, they might each have the goal 'the sunbathers on this beach (whoever they are) do not sit right next to each other'. This goal might be out in the open between them, for example if each sees the others act in ways that promote this goal, or simply if they live in a society where most people have this goal when at the beach. So, (1) each endorses the same goal, and it's out in the open between them that the goal is endorsed by each. They are also (2) disposed to act responsively to one another (insofar as they encounter one another) as they work towards that goal: simply, they are disposed to put their towels at a respectful distance from each other. And they also (3) lack group-level procedures for forming decisions that are attributable to the group as such. Such a procedure might easily come into existence—as it does in the group of friends who must decide where to go for lunch together—but until the members establish a procedure in the way the group of friends did, decision-making is going on only at the level of the members. I'll argue for this in Section 2.3.2.

1.3.3 Combinations

A combination is constituted by agents who do not together constitute a coalition or a collective. Examples of combinations include 'men' (since common advantage does not suffice for a common goal), 'humanity', 'the international community', 'the people in this pub', and 'me, you, and Shakespeare'.

From this, we can see that some combinations have their members essentially, while others don't. The combination of 'me, you, and Shakespeare' can't remain the same group while its membership changes: if we subtract me from the group, then we get a different group altogether. By contrast,

other combinations do not have their members essentially. 'The people in this pub' is like this. It refers to whoever happens to currently be in the pub. If the guy with the pram leaves, then that is a change to the group 'the people in this pub', but it is a change that has occurred to an entity that persists once he leaves. Thus the notion of a combination encompasses what we might think of as mere 'sets' or 'collections' (that change their identity every time they change membership), as well as socially significant groups (that do not change their identity every time they change membership, though they might with significant changes in membership). (Pettit (2014, 97) thus distinguishes 'sets' or 'collections' from 'groups'; likewise French (1984, 5) uses this criterion to distinguish 'aggregates' from other types of group.)

Obviously, some combinations are far more important than others for moral, social, and political philosophy. Combinations such as 'men' or 'white people' or 'the working class' are important philosophical objects. But I treat all these varied groups together because—no surprises here—they share the same status as duty-bearers. This will become clear as the discussion proceeds in Chapters 2 and 3. Let me forestall potential scepticism by noting that an agent's membership in a combination can be a crucial source of duties for them, and that the extent to which combination-membership plays this duty-generating role varies hugely amongst different combinations. Thus, the fact that an individual is wealthy, or white, almost certainly generates duties for them as an individual that they would not have if they were not wealthy or white. Such duties will not be my concern in this book. Instead, I am concerned with whether the group per se can have duties and, given that it does or doesn't, what this fact alone implies for the duties of a member. Some facts about a combination—such as the fact that its members are statistically likely to be advantaged or disadvantaged—produce duties for the members of that combination which non-members do not have. But these will not be my concern, since they do not derive from the fact that the group itself does (or does not) have a duty. This caveat brings us to the question of what the Tripartite Model does, and doesn't, try to capture.

1.4 The Scope of the Model

The Tripartite Model is a general framework for conceptualizing both (i) group duties and (ii) members' duties that derive from a group's duty (or its absence). It tells us which groups are candidates for duties and what this implies about the basic structure of members' duties.

I'll have more to say about groups than about duties. However, to clarify the subject of the Tripartite Model, I should say something about duties.[5] I understand duties as objective moral requirements to perform certain actions or produce certain outcomes. Requirements are presumptively decisive in decision-making, even though that presumption can be overridden. For example, one duty can override another in a context where both cannot be followed. Nonetheless, it is in nature of duties that the duty-bearer should view both duties as having a presumptive claim to compliance. What's more, any overridden duty leaves behind a moral remainder: it is appropriate for the duty-bearer to feel regret or remorse when failing to follow an overridden duty.

5 I thank Christian Barry, Nic Southwood, David Killoren, and Marion Smiley for discussion of this.

I am focused on moral duties—not legal, epistemic, rational, social, conventional, prudential, or other kinds of duties. This makes the Tripartite Model useful to normative moral and political philosophers—even if it diminishes the model's usefulness to lawyers or sociologists. Another restriction is that I'm concerned with duties to perform actions or produce outcomes. I'm not concerned with duties to have certain emotions, attitudes, or dispositions. The latter are trickier subjects for a theory of duties, because perhaps no agent has immediate voluntary control over them. So I leave them for another day. (I tackle organizations' duties over emotions in Collins (2018).)

Many of the duties with which I'll be concerned are grounded in the value of outcomes, where an 'outcome' is roughly the same thing as a 'state of affairs', 'situation', 'scenario', or 'region of possibility space'. My examples tend to build upon the idea that—at least sometimes—an entity can have a duty to perform some action simply because that action would produce a highly valuable outcome. If the entity has a good shot at producing the outcome, we might frame this by saying the entity has a duty to produce the outcome. At other times, if the entity's influence over the outcome is more tenuous, we might say the entity has a duty to perform actions that increase the likelihood of the outcome or a duty to try to produce the outcome. Chapter 4 will have more to say about that distinction. The point here is that, in both cases, the duty is grounded in the outcome's value.

To be clear, I mean for the Tripartite Framework to be applicable to all duties—whether outcome-based or not. In Chapter 7, for example, I will use the example of a collective having a duty to ensure its suppliers are paid fair wages. This is best understood as an agent-relative duty not to exploit or be complicit in exploitation—not as a duty to produce some agent-neutral outcome. More generally, I assume that collectives can have the full range of duties-to-act that individuals can have, including duties to keep promises, avoid doing harm, return benefits received from injustices, be partial to special relatives, and so on (as argued in Collins and Lawford-Smith 2016). At other times, however—particularly when discussing coordination duties in Chapter 4—the examples will be cleaner if we take an outcome-based approach. And I will ground coordination duties, in particular, in the value of outcomes.

This does not take sides on first-order normative ethical theorizing. One needn't be a full-blown consequentialist to believe that at least some duties (for example, the coordination duties that Chapter 4 posits) are grounded in the value of outcomes. Any moral theory that has a concern for consequences can agree. This will include almost all forms of deontology, as well as forms of virtue ethics that countenance the notion of 'duty'. I do not mean to suggest that all moral duties—including collectives' duties and membership duties—are grounded solely in the value of the outcomes that would be produced if those duties were followed.

Another clarification is that I am concerned with objective duties—not subjective, evidence-relative, or other kinds of 'relativized' duties. My concern with objective duties derives from the fact that objective duties seem to analytically ground relativized duties. For example, one's subjective duty is (roughly) what one's objective duty would be if one's beliefs were true (it is one's duty relativized to one's beliefs); one's evidence-relative duty is (roughly) what one's objective duty would be if one's evidence reflected all the relevant facts (it is one's duty relativized

to one's evidence); and so on. Objective duties have a privileged status in that (i) relativized duties can be analysed in terms of objective duties and (ii) objective duties are not relativized to anything—or, rather, they are relativized to the totality of all truths.

Objective moral duties' privileged status is reflected in the following: if you know what someone else's objective moral duty is, then it is permissible for you to interfere with them (potentially exposing them to cost) to make it the case that they do their duty. For example, if I'm about to put poison into my friend's drink, but I reasonably believe the poison is gin, then I am about to violate my objective duty not to poison my friend. However, my subjective (belief-relative) and evidence-relative duties are not about to be violated—because I do not believe, and my evidence does not indicate, that the liquid is poison. It is permissible for you to violently knock the poison out of my hands. Such a permission would not follow from you knowing my subjective or evidence-relative moral duty—but it does follow from you knowing my objective moral duty.

Objective moral duties are also privileged in that their violation leads to a moral remainder (as intimated above). Their violation produces a remainder even when it doesn't produce blameworthiness. For example, if I put poison into my friend's glass while reasonably believing the poison is gin, then it is appropriate for me to apologize, or seek to provide recompense, or feel agent-regret (of the kind discussed in Williams 1985). These are appropriate responses to the moral remainder, even though I'm not blameworthy. Of course, these responses (and more) are also appropriate if I believe the poison to be poison, or if my evidence suggests it's poison. But objective moral duties are special in that their violation produces a moral remainder without producing blameworthiness. (The possibility of duty-violation without blameworthiness is one reason to be sceptical that duty and blameworthiness are two sides of the same coin.)

Objective moral duties are further special in that they make certain attitudes (more) fitting. For example, suppose one's objective, subjective, and evidence-relative duties align—that is, one believes one ought to do something, and one's evidence suggests one should do that thing, and one objectively ought to do that thing. And suppose that one does the thing, because one believes one ought to do it, where one believes that because one's evidence suggests it, where one's evidence suggests it because it is, objectively, one's duty. Under these sorts of conditions, it is appropriate for others to bestow praise, commendation, approval, endorsement, or at least condoning. Of course, these reactions would also be appropriate if one did merely one's subjective or evidence-relative duty—but these reactions are more appropriate, and are appropriate in larger quantities, when we add that what one has done is one's objective duty.

Thus, objective moral duties are important in three ways: they generate permissions for third parties to interfere to ensure compliance; they generate a moral remainder without (necessarily) blameworthiness when they are violated; and they render certain attitudes (more) fitting when they are done—complied with—in the right way (Section 3.3 will have more to say about doing a duty 'in the right way'). Thus, a focus on objective moral duties is appropriate. I assume that everything I'll say about the objective moral duties of groups and their members can be appropriately tweaked to produce accounts of groups' and members' subjective, evidence-relative, etc, moral duties. But I will remain focused on objective moral duties.

Objective moral duties are weightier than mere objective moral reasons. Reasons, I assume, are cheap: I have some reason to eat your car, insofar as it contains iron, which is good for me. While reasons are mere considerations in favour of some conduct, duties demand—require—that conduct. As I said above, duties are presumptively decisive in decision-making: their bearers should presume that the duty is not overridden, outweighed, or undermined—and only in special circumstances should that presumption be overridden. I am neutral on whether this book's arguments have straightforward analogues when it comes to groups' reasons.

As part of my neutrality on first-order normative theorizing, the Tripartite Model doesn't answer the substantive question of which specific duties are held by which specific groups, and which specific duties members have on this basis. Answering that question would require a full-blown first-order moral theory. The Tripartite Model must be combined with such a normative ethical or political theory if we want to know which specific group duties exist and what specific duties they imply for members. If you take (1) your favourite moral theory, plus (2) empirical facts about the structure of each group, then (3) plug all that into the Tripartite Model, you'll get (4) substantive conclusions about which specific groups have which specific duties, and what that implies for each group's members.

To return to an example with which I began: the Tripartite Model tells us that the UK is the kind of group that can bear duties, since the UK is a collective. The UK's being a collective follows from empirical facts about the UK, plus the definition of 'collective' provided by the Tripartite Model. So, it makes sense to say that the UK—in its own right—has a duty to provide safe harbour to people fleeing conflicts. But to establish that the UK really does have that duty, we would have to provide a substantive moral justification for that particular duty (where that justification is, most likely, based on the value of living free from conflict). This requires moral theorizing of the kind this book leaves to one side.

Likewise, the Tripartite Model tells us that the group 'conservatives' and the group 'the international community' are not the kinds of groups that can bear duties. Again, this follows from the empirical facts about these two groups, plus the Tripartite Model's verdict on groups with that kind of structure. So, it is a category error to say that these groups bear duties to defend liberal-democratic civilization or to bring peace to Syria. But the Tripartite Model offers an alternative interpretation of these duty attributions: a duty attributed to conservatives makes sense if it's conceptualized as a conjunction of duties, each of which is held by one conservative, each of which is a duty for that conservative to take steps—based on coalition-reasoning with other conservatives—towards the defence of liberal-democratic civilization. And Ban Ki-moon's statement makes sense if it's understood as saying that each (individual or collective) member of the international community has a duty to take steps—based on I-reasoning—towards bringing peace to Syria. Again, substantive normative argument (based on the value of liberal-democratic civilization or the value of peace) would be required to establish that the members of these groups really have these coordination duties. This is why my conclusions are limited to what 'makes sense'.

This might seem like a modest result. But since the Tripartite Model applies to all groups—of all conceivable shapes and sizes—this result is significant. All you have to do is 'plug in' your

favourite moral theory and some facts about the groups you're interested in, and the Tripartite Model will spit out an abundance of coordination duties, collectives' duties, and membership duties. The book's conclusion is that collectives can bear duties, that coalitions and combinations cannot, and that these twin results respectively imply membership duties and coordination duties. This constitutes a large step towards a complete ethics for groups.