1 *Introduction*

International political theorists have become increasingly interested in questions about collective moral agency and responsibility (Erskine 2003b), the ethics of global governance (Franceschet 2009), the relationship between sovereignty and human rights (Rawls 1999a; Wheeler 2000; Brown 2002; Beitz 2009), and the philosophy of international law (Besson and Tasioulas 2010). At the same time, moral and political philosophers have begun to look systematically at the relationship between human rights, on the one hand, and human rights duties, on the other (Shue 1996; O'Neill 2001; Pogge 2007b). In International Relations (IR) scholarship, questions about the roles and capacities of non-state actors – though not entirely new – have taken on a new dimension with a shift to debates about non-state authority in world politics (Hall and Biersteker 2002; Cutler 2003; Avant *et al.* 2010; Abrahamsen and Williams 2011). And in legal theory, there is an active theoretical discussion of the transnational human rights duties of states (Langford *et al.* 2013). Despite significant interest in these topics, there has still been insufficient scholarly analysis of the responsibility for human rights of non-state actors that is grounded in international theory as its main disciplinary starting-point.[1] This book makes a significant contribution by providing an original theoretical account of the responsibility for human rights of non-state actors, using transnational corporations (TNCs) as the main example for the argument.

It is particularly important for a book of this kind to be considered now, because of the rapid development of a policy agenda that aims to use international and transnational legal mechanisms to assign such responsibilities to those actors (Ruggie 2013). This agenda takes a variety of forms, but it has most notably manifested in recent years in John Ruggie's work since 2005 on business and human rights at the UN Human Rights Council (Gibney and Emerick 1996; Teitel 2005; Ruggie 2008b; Jerbi 2009; Giannini and Farbstein 2010). At first glance, it may seem as though this kind of policy project makes perfect ethical sense. Yahoo's Chinese subsidiary, for example, has been

[1] For excellent international-law analyses, see Ratner 2001; Alston 2005; Clapham 2006.

accused of contributing to forced labour and torture, because it complied with an official request from the Chinese authorities to release the IP address and the personal details of activist journalist Shi Tao to the Chinese police (BBC News 2007; see also Miller 2009). Subsequent events in the legal and political environment led rival Google to threaten, in 2010, to close its Chinese subsidiary altogether. Royal Dutch Shell has also been the subject of ongoing lawsuits since the 1990s, when, victims claim, it contributed to the detention and eventual execution of human rights and environmental activists in the Niger Delta region of Nigeria (BBC News 2009a; Center for Constitutional Rights 2009; Center for Constitutional Rights 2012). Surely holding these companies to account as human rights violators must be a good thing; so the argument goes. But once one probes more deeply, it becomes clear that the policy project is a legal and political innovation that does not conform to any standard account of what it means, fundamentally, for any agent to have a distinct kind of responsibility linked specifically to human rights rather than to other global concepts or practices.

The argument and layout of the book

There are two very broad accounts of which actors in world politics have responsibility for human rights. The first broad account, which can be called the 'universalist' account, is that human rights responsibilities are those that we all have to each other by virtue of our common humanity. On this view, any agent can be a human rights violator when he or she acts, or fails to act, towards others in a way that respects this common humanity. In other words, anyone who causes harms of a particular kind to another human being is a human rights violator, according to this first account. The second broad account, which can be called the 'state-centric' account, is that human rights responsibilities fall only on states. The reasons offered for this vary. Legal-positivists typically believe that human rights duties are legal duties that need to be agreed, in international law, by states, because of the constitutive rules of international-legal validity. Political theorists and philosophers offer a different set of reasons, including that states have unique capacities that other actors lack and/or that states are the sole legitimate political authorities in contemporary world politics. These facts, in turn, can be linked theoretically to the normative need for a state-centric approach to human rights responsibility and accountability.

To explain the difference between the two accounts, consider this example. If a neighbour terrifies a person to the point where that person is afraid to leave her home, then the neighbour might have done something criminally wrong, but is he a human rights violator? According to the universalist account he is, because he has violated the neighbour's security and effectively confined her to her home. But according to the state-centric account, he is not a human rights

violator, because all of these actions occurred within the private sphere: they did not involve a state actor. By contrast, if a police officer had done the same thing, or had not responded to a call for help, or if the government had not ensured that systems were in place such that help could be provided when needed, then a human rights violation might very well have occurred, according to the state-centric account.

The point, and the problem that this book is designed to address, is that the project of identifying TNCs as specific kinds of human-rights-responsible agents problematises *both* of these standard theoretical accounts of what human rights responsibility is, which kind of agent can bear it, and why. Employees of private security companies, for example – especially in mixed governance zones such as the Niger Delta – do effectively confine people to their homes. If all moral agents can be human rights violators, and if corporations can be viewed as moral agents, or at the very least as collections of moral agents, then there would be no reason to separate these companies out from other agents, as unique bearers of human rights responsibility. The project would be normatively justified, but in a trivial way: in a way that would fail to make sense of the impetus for a specific policy agenda dedicated to human rights and global business. On the other hand, if only states can bear human rights responsibility, then to add companies into the mix would count as a radical change to the international human rights regime, and this, in turn, would require thought and reflection about the extent to which this seemingly major break with existing human rights practice would be justified.

Digging beneath the surface of those two very broad accounts, there are four narrower candidate theories of responsibility for human rights, all of which will be assessed by this book. The first theory is the 'legalistic' theory. This theory says that states and international institutions can simply create new international law that assigns legal human rights obligations to transnational corporations, or can provide interpretations of existing international law showing that TNCs already have such obligations; if they do so successfully, then TNCs straightforwardly have human rights responsibilities, and these can legitimately be enforced in any domestic, regional and/or international jurisdiction. The second theory is the 'universalist' theory, which says that all moral agents have human rights responsibility, and that TNCs therefore simply have the human rights responsibilities that we all have, provided that they are moral agents. The third theory is the 'capacity approach' to human rights responsibility. According to the capacity approach, as I define it, whichever agent has the capacity most effectively to protect and to provide for human rights has a specific responsibility to do so. All moral agents might have latent human rights responsibilities, but – distinguishing the capacity approach from the universalist approach – they only become activated once a particular agent in a particular set of circumstances is identified as having greater capacity than others to protect

or to provide for human rights. Fourthly and finally is the 'publicness' approach to responsibility for human rights. According to the publicness approach, agents that are relevantly public, but not agents that are relevantly private, have specific human-rights-based obligations. The meaning of the approach rests to a great extent on how public and private are defined for this purpose. To be public enough for human rights does not imply being public enough for other purposes and practices.

The book makes the following argument. Chapter 2 lays the groundwork for the relevance of the book by providing an overview of how and why 'TNCs and human rights' is an issue that has come onto today's international policy agenda. It introduces a small number of paradigmatic contemporary examples of alleged human rights violations by corporations, each of which captures important dynamics that are reflected in a broader range of cases. Chapter 3 then challenges the validity of an analysis that focuses exclusively on the question of what the law is. It explores the idea of a 'delinquent state', in which a state's actions and decisions are the primary drivers of candidate human rights violations with which companies are then accused of complicity (see also Erskine 2010). A delinquent state might require non-state actors to contribute to harms to individuals that are *prima facie* examples of human rights violations. When the focus is placed on TNCs that operate in contexts which resemble the 'delinquent state' model, it becomes clear that normative and political questions about the respective authority of sovereign states' rules on the one hand, and international-legal rules about human rights on the other, need to be addressed much more systematically from an extra-legal theoretical perspective.

Chapter 4 begins to do this in detail. It assesses the idea that responsibility for human rights is universal: that all moral agents have it. The chapter advances contemporary discussions about the human rights responsibilities of non-state actors by distinguishing three kinds of responsibility that are linked to human rights practice: (a) responsibility to protect and to provide for human rights, (b) responsibility to refrain from harming human rights and (c) responsibility to respect human rights.[2] Ruggie's (2008b; 2011) project at the UN Human Rights Council uses language that, on the surface, is similar. The UN Guiding Principles on Business and Human Rights, which developed under his leadership, revolve around the central idea that states have the duty to protect human rights, and that corporations have responsibilities to respect human rights, where the latter category is understood to include: responsibility to refrain from the active violation of basic rights; responsibility to conduct

[2] This trichotomy resembles Shue's (1996, 35–64) distinction between duties to avoid violating rights, duties to protect rights and duties to aid victims of rights violations. It will become clear in Chapter 4 of this book that my account is somewhat different.

due diligence before investment in a context where human rights violations might occur; and transparent reporting on human rights operational impacts. In fact, the framework incorrectly collapses two very different kinds of responsibility, responsibility to refrain from harming human rights and responsibility to respect human rights, into one. Responsibility not to harm human rights is non-discretionary, and requires agents to possess only a basic threshold of moral agency. Responsibility to respect human rights is discretionary, and requires agents to construct (in a sound way) and to refer to (in a valid way) thicker frameworks of ethical judgement. It is true that all moral agents have responsibilities both to respect human rights and also to refrain from harming human rights. However, it is not true that all moral agents have responsibility to protect and to provide for human rights. A historical interpretation of the point and purpose of human rights practice will show that – although all three of these kinds of responsibility are a part of human rights practice – responsibility to protect and to provide for human rights, which do not fall on all moral agents, but rather, fall on specific agents for specific reasons, constitute the core of responsibility for human rights. Some, but not all, of the responsibilities linked to human rights practice are universal. The non-universal kind of responsibility should be the primary focus of theoretical questions about who bears human rights responsibility, what this means, and why. This structural account of responsibility for human rights does not immediately and necessarily imply that states and only states are the kinds of agents that should have specific responsibility in international law and in political practice. Further analysis of who has specific responsibility, and why, is needed.

Chapter 5 assesses the capacity approach to human rights, which generates a method that can be used to identify which agent or agents bear specific human rights responsibility in specific circumstances. The capacity approach has at least three strengths. Firstly, it offers a method to assign human rights responsibility to specific agents. Secondly, it provides a principled middle way between universalist and state-centric approaches to responsibility for human rights, because states often, but not always, have the capacity effectively to protect and to provide for rights. Thirdly, it has both an ideal-theory mode, in which it asks who has human rights responsibility, and a non-ideal-theory mode, in which it asks who else this responsibility might be assigned to (if anyone) if the 'ideal' responsibility-bearer is unwilling or unable to act on his, her or its responsibilities.[3] This is particularly important, because human rights problems

[3] For a good review of the literature on the distinction between ideal and non-ideal theory, see Valentini (2009). I take Buchanan's (2004, 55) interpretation of Rawls as standard: 'Ideal theory sets the ultimate moral targets, articulating the principles that a just society or a just international order would satisfy, on the assumption that there will be full compliance with these principles. Nonideal theory provides principled guidance for how to cope with the problems of noncompliance and how we are to move closer toward full

become most apparent when those who should have human rights responsibility act irresponsibly. However, the capacity approach suffers from two major, related, deficiencies. It focuses entirely on obviating costs to potential victims, and it does not offer an account that stands up to critical scrutiny of how to take the burdens on potential responsibility-bearers into account. I shall analyse this through a distinction between agent-centred ('type-1') and victim-centred ('type-2') errors (Barry 2005). A type-1 error would occur if the human rights regime assigned human rights responsibility to an agent unjustifiably, whereas a type-2 error would occur if responsibility were *not* assigned to a particular agent, leading to an unjustifiable cost to a human rights victim. The capacity approach also treats 'responsibilities to others' as a monolithic kind of practical ethics. A more nuanced approach to responsibility for human rights needs to be able to distinguish between human rights responsibility, on the one hand, and different responsibility practices – for example, criminal responsibility, remedial responsibility, blame responsibility – on the other, at the stage of assigning potentially burdensome responsibilities to particular agents. Because of these two deficiencies, the capacity approach should not be adopted as a stand-alone approach to responsibility for human rights. These three strengths and two weaknesses represent a set of five criteria, all of which are necessary for a theory of human rights responsibility to meet in order to be sound.

Chapter 6 defines and assesses the publicness approach to responsibility for human rights. It argues that the publicness approach is able to match the strengths of the capacity approach while overcoming its deficiencies. There are five possible ways to identify specific agents as 'public' for the sake of assigning human rights responsibility: (1) as agents who provide empirical collective goods, such as security, roads, hospitals and schools; (2) as agents that provide membership in a political community and take decisions about 'insider' status; (3) as agents that provide political responsiveness for the members of a political community; (4) as socially constructed agents, offices and roles (whereas individuals, separated from all their roles, are relevantly private); (5) as agents with a political role, who exercise *de facto* (not necessarily legitimate or normative) political authority. These versions overlap in significant places, and are ultimately used in combination rather than treated as discrete and competing definitions of 'publicness'. The publicness approach is the best way to analyse who is specifically responsible for human rights.

Chapter 7 concludes the book by synthesising the argument of its substantive chapters, and by clarifying the argument's implications for theory and

compliance with the principles of ideal theory.' When I discuss non-ideal theory in this book, I focus on how to cope theoretically with the possibility (even likelihood) of non-compliance with ideal theory rather than on pragmatic questions about what to do to change the behaviour of actors who are identified as responsible in ideal theory but who fail to act in accordance with it.

for practice. It re-states the book's overall endorsement of the publicness approach, while diffusing one of the main objections that can be raised against it: that it might permit the under-fulfilment of human rights outcomes in the short term if no relevantly public agent can be identified and/or if the only public agents available act irresponsibly. The chapter also suggests the way the argument can be adapted and applied in the future, in further research about the responsibilities of a broader range of contemporary non-state actors.

Human rights and human rights responsibility

Three clarifications are important before proceeding to the main argument. Firstly, there are two views in relation to which I position my own view about the relationship between 'human rights' and 'human rights responsibility'. Questions about this relationship are sometimes referred to as questions about the co-relativity of human rights and human rights responsibilities. The first view in relation to which I position my own view is the idea that a full-fledged theory of human rights needs only to be able to generate and to justify a list of important objects (material or abstract), such as security, political participation, health and education, to which all humans are morally entitled (for example, Griffin 2008). A moral defence of the idea (a) that all individuals ought not to live in poverty does not get one automatically to the conclusion (b) that some agent or agency ought to be responsible for alleviating poverty. To get from (a) to (b), one would need to provide a sound argument fleshing out the intermediate steps. An ability or inability to provide an argument that gets from (a) to (b) has no independent bearing on whether or not (a) is a true or reasonable premise. The second view in relation to which I position my own view is O'Neill's (2000, 97–111; 2001, 185) idea that discussions of human rights supervene entirely on discussions about duties. She aims to re-frame discussions about rights around discussions about duties, arguing that discussions about rights are 'indeterminate and ineffective' unless those rights can be shown to be grounded by duties that attach to certain agents.

Discussions about the material and abstract objects that are of fundamental moral importance to individuals, and discussions about moral responsibilities, are important ethical discussions in and of themselves. But neither, by itself, is a full-fledged human rights discussion. An adequate theory of human rights needs not one discussion, but two separate and separable discussions: a discussion of objects to which all humans are morally entitled (this can include a discussion about whether there even are any such objects), and another discussion about who, if anybody, has which kinds of responsibilities as a result of all individuals having at least some important moral rights. Both are important ethical and practical discussions. But neither, by itself, is a full-fledged discussion of human rights.

Full-fledged human rights discussions occur when one asks whether and how the conclusions of each of these two discussions fit together. If one can successfully argue that freedom from poverty is an important normative right (defined for preliminary purposes as a right that all humans ought to have even if they do not currently have it in positive terms), but if one fails to show how and why important normative rights generate 'human rights responsibilities', then one has not produced a full-fledged theory of poverty as a human right. By the same token, if one can successfully argue that the moral responsibility to alleviate poverty exists, and ought to generate specific duties that fall on agents or agencies, but if one cannot successfully argue that all individuals have a normative right not to live in poverty, then one also has not produced a full-fledged theory of poverty as a human right. The two pieces need to be in place individually, and once in place, they need to meet a further test of fit.

This book assumes, for the most part, that the first piece can be put in place successfully. I take it as uncontroversial that a good moral philosopher can defend the importance of objects such as security, political participation, education, health and so on to individual persons. In what follows, I shall therefore focus on the second discussion, the discussion of human rights responsibility, and also on the 'further' or 'full-fledged' discussion of how, if at all, the second discussion links up with the first.

Can transnational corporations be moral agents?

The second clarification revolves around the question of what moral agency is, and whether transnational corporations count as moral agents in contemporary international politics. The answer to this question in what follows is just an introduction; it will help readers who are interested in this issue to tie together the details that are provided in Chapters 3, 4 and 6.

In her own contribution to an excellent volume that examines the question of whether corporate and institutional actors are moral agents that can have responsibilities in international relations, Erskine (2003a, 23–24), drawing from French (1984), claims that in order to count as a moral agent, a corporate or collective actor needs to have, at a minimum, a decision-making structure that entails a 'capacity to access and process information', as well as a degree of decision-making unity – what I would call a single self – that persists over time. French (1984) uses the language of an 'identity' that persists over time. However, the term 'self' is a more accurate one, which brings the view that he expresses more closely in line with the way that social psychologists understand the difference between 'self' and 'identity'. Identities include 'black', 'white', 'man', 'woman', 'gay', 'straight' and so on, whereas the 'self' is the container inside of which various individual identities co-exist (Stryker 1980). There is a further distinction still that can be made between identities and roles, which will become important as the book proceeds. Individuals often

feel that they inherently *are* their identities (I will not bore readers with the *Oxford English Dictionary* definition of identity). This fact can range from empowering, for example, gay people can justify political claims on the basis that they cannot change who they are; to problematic, for example, there are aspects of socially/culturally constructed gender identity that can constrain female autonomy. By contrast, roles can more easily be slipped in and out of: the police officer who takes recreational drugs on the weekend while off duty, the father who behaves differently at home than when out at the pub with his friends, the bohemian rock star who eventually becomes a philanthropist or a lawyer. The relationship between agency, roles and responsibility for the sake of human rights is one of the threads that runs throughout this book.

In order to meet the first criterion for agency, transnational corporations would need to be able to process, interpret and act on rules. The ability to interpret – rather than simply the ability to process and to act – is important (DeWinter 2003, 139). A robot could process an instruction to make a cup of coffee or to clean the floor, but this does not make it a moral agent. By contrast, no robot, of which I am aware, would be able to act successfully on an instruction to act such that individuals are protected from violations of due process rights. The latter requires understanding and interpretation as to what counts as a violation, which in turn requires a conscious, human judgement. For this reason, corporate and collective agents require a sentient element within them. However, they also need to be more than the sum of their parts in order to count as corporate. Due to corporate constitutions or institutional structures, which dictate how individual (human) actors within corporations must take decisions – and which prescribe particular factors that are (and are not) to be taken into account before arriving at a final decision – the agency of corporations, if indeed they *are* corporate agents, cannot be reduced purely to the agency of individual members. The positive argument that corporations meet this first criterion is well established (Pettit 2007). It is worth pointing out that companies meet this criterion even more closely than most other candidates for corporate agency, such as nations, international organisations, transnational identity groups or genocidal mobs.

The question of whether transnational corporations meet the second criterion, a unitary self, is more complicated and controversial. The head offices and the subsidiaries of corporations, which are incorporated in separate countries or jurisdictions, are treated by the law as separate persons – separate decision-making agents – and not as a unitary whole. This legal *status quo* seems to clash with the facts of the matter about corporate decision-making, which, in many cases, involves the various parts of transnational companies working together explicitly in what they view as a common enterprise. These facts about corporate decision-making do matter. *Some* transnational corporations might be best characterised as several corporate agents rather than as a single corporate agent. Others are best characterised as single agents. This depends

entirely on the extent to which there is, or is not, actually a single decision-making structure that connects the units and enables them to work together for a common purpose. The specific companies that I shall discuss in this book, for example, Yahoo and Shell, do have a considerable degree of decision-making unity, in practice though not necessarily in law. Indeed, one of the aims of Chapter 3 of this book is to discuss the normative problems that arise precisely because a single agent is treated by the rules of jurisdiction as a set of legally separate companies.

It is not enough just to say that transnational corporations are unitary actors that have the structures and capacities needed to take decisions, and that they therefore count as moral agents for the sake of assigning responsibility for human rights to them. The kind of decision-making in which they need to be able to engage matters as well. One of the arguments that this book will advance is that different kinds of responsibility practices require responsibility-bearers to exercise different degrees of judgement. The two criteria discussed above represent a minimum threshold for corporate moral agency. But beyond this, there are higher thresholds. The threshold that an actor needs to meet in order to count as a responsibility-bearer within a specific responsibility practice (for example, human rights practice) is dependent, at least in part, on the contours of that particular form of responsibility. In Chapters 3 and 4, I unpack this claim. I shall distinguish between non-discretionary and discretionary duties. The former require agents to identify, to understand and to follow the relevant rules. The latter require agents to be able to formulate and to work with what philosopher Charles Taylor (1985) calls 'strong evaluations', which, on my interpretation, means: values that are actually valuable (rather than trivial), which an agent authentically holds and has incorporated into his or her self. To be responsible for acting on non-discretionary duties, agents need to meet a minimal threshold of moral agency. To be responsible for acting on discretionary duties, agents need to meet a more robust conception of moral agency, which requires the development of a legitimate set of principles and values, and the ability to refer to those, in a valid way, when acting.

I argue in Chapter 4 that in order to be responsible (a) not to harm human rights and (b) to protect and to provide for human rights, corporations would need to meet the more minimal threshold. In order to be responsible (c) to respect human rights, corporations would need to meet the more robust threshold. I shall furthermore offer an account of whether and why corporations might be thought to meet *both* of these thresholds for moral agency. That companies can meet the first threshold is not particularly controversial. There are two ways that companies could meet the second. Firstly, companies can be constitutionally and/or institutionally set up so they have a body of values and principles, including, but going beyond, responsibilities to shareholders, to which they must refer in taking operational decisions. Secondly,