



PROJECT MUSE®

The Moral Nexus

R. Jay Wallace


Published by Princeton University Press

Wallace, Jay.

The Moral Nexus.

Princeton: Princeton University Press, 2019.

Project MUSE., <https://muse.jhu.edu/>.



➔ For additional information about this book
<https://muse.jhu.edu/book/63235>

2

The Problem of Moral Obligation

IN THIS CHAPTER, I look at the hoary issue of the normative significance of moral requirements in the first-person perspective of deliberation. Moral conclusions are customarily treated as considerations that matter within an agent's practical decision-making. That a course of action would be impermissible, for instance, or morally the right thing to do, are conclusions that appear to have direct relevance for practical deliberation, which agents who are reasoning correctly will take appropriately into account in planning their future activities. The philosophical problem in this area is accordingly often understood to be the problem of making sense of the reason-giving force of morality. That is, an account of moral rightness or permissibility should shed light on the standing of these considerations as reasons for action, which count for and against actions in the first-person perspective of agency.

In my view, however, this conventional understanding seriously underdescribes the challenge that faces a philosophical account of morality. The problem is not merely to explain why moral conclusions have some normative significance in the perspective of practical deliberation. It is to explain their apparent status as obligations, considerations that function in a very distinctive way within the practical reflection of agents who take them properly into account. This is the problem that was raised, very forcefully, in G. E. M. Anscombe's justly famous paper on "Modern Moral Philosophy."¹ Anscombe highlighted the fact that morality, on the modern conception of it, appears to involve a class of considerations that make normative demands on the agents to whom they apply. We feel that we are morally bound or obliged to act in accordance with our thinking about what it is morally right or impermissible to do. But Anscombe maintained that this conception of moral obligation is

philosophically unintelligible. We could perhaps make sense of moral obligation within the framework of a “law” conception of ethics, which holds that moral requirements result from the demands of a beneficent and omnipotent deity. But when we detach morality from this kind of theological framework, as we must if we are to understand it as a secular phenomenon, the sense of obligation proves fundamentally elusive. Anscombe concludes that we should dispense altogether with the special notion of moral obligation, and find a way of doing ethics without this distinctive concept.

Anscombe seems to me too quick to give up on the notion of moral obligation, which is central to our ordinary understanding of the significance that morality has for individual deliberation. But she is right that this notion is vulnerable if we are unable to provide a philosophical account that makes sense of it in secular terms. My aim in this chapter is to show that the relational approach to morality is well suited to meet this important challenge. Understanding moral obligations in relational terms, as duties that are owed to other parties, renders them intelligible as considerations that have the distinctive normative force that Anscombe found elusive. Indeed, the relational approach offers the most plausible way of understanding interpersonal morality as a coherent and unified set of obligations, and this is among the most important theoretical advantages of interpreting the moral in fundamentally relational terms.

I begin my discussion by setting out, in section 2.1, the general problem of making sense of obligation as a deliberative phenomenon, and sketching some of the philosophical approaches available for addressing this problem. In section 2.2, I focus specifically on moral obligation, identifying some additional features of it that an adequate moral theory needs to accommodate; I also canvas some familiar moral theories, highlighting the problems that those theories inevitably face when it comes to making sense of moral obligation as a deliberative phenomenon. In section 2.3, I present the relational account as a superior approach to understanding the key features of moral obligation. Duties that are owed to another party are often understood to exemplify the original notion of an obligation, and it makes sense that we should feel ourselves bound by moral considerations if they represent directed obligations in this very particular sense. In section 2.4, I explore some implications of the relational account of obligation, contrasting it with the voluntarist approach that Anscombe and many other philosophers take to be the most natural alternative.

2.1. Practical Requirements: The Basic Challenge

Moral considerations, or at least a central class of them, familiarly exhibit what I have elsewhere called deontic structure.² That is, we commonly take morality to consist in large part of a set of basic requirements on the will, which preclude our doing some things and demand that we do others. Morality is not merely a set of considerations that count in favor of or against doing certain things, but a source of obligations, which are practical requirements that function as fixed constraints on what we may or may not do.

Practical requirements of this kind are highly distinctive; they differ intuitively from the reasons proper to many other normative domains, such as those of professional or personal advantage.³ The fact that it would be entertaining for me to attend the concert next week at the Philharmonie is a reason for trying to get tickets, perhaps even a weighty or a compelling reason. But it isn't one that I take into account in reflection as a demand or an obligation. By contrast, if I've promised you that I would get us tickets to the concert, then the case acquires a moral complexion that shifts it into a different normative gear. In particular, it now seems not merely that there is something that speaks in favor of my endeavoring to get us tickets, but that I am subject to a presumptive requirement to do so.

It is something of a commonplace these days to understand normativity in terms of reasons, and to interpret reasons in turn as considerations that count in favor of the attitudes and actions that they support.⁴ A reason for going to the supermarket on the way home from work is a consideration that might be set on the positive side of a notional ledger, and weighed against considerations entered on the other side, as counting against the action. But not all normative considerations are correctly understood in these terms. In particular, practical requirements seem to function very differently within deliberation than the kind of pro or contra considerations that might be weighed against each other in these ways. As Samuel Scheffler has written, such requirements intuitively function as "presumptively decisive" reasons for action and response.⁵ They are not considerations that are ordinarily weighed against other, potentially competing reasons for action, but operate rather as exclusionary reasons (in Joseph Raz's influential phrase), which defeasibly block the normative force of considerations that in other contexts would serve as perfectly respectable reasons for action and response.⁶

Semantically, this dimension of practical requirements is reflected in the fact that they characteristically find expression in claims about what an agent

must do, rather than about what the agent ought to do. “Must” and related deontic expressions (such as “have to”) signal the peremptory or decisive aspect that distinguishes practical requirements from other kinds of normative consideration. A further dimension of the contrast between ordinary “counting in favor of” reasons and practical requirements has to do with the idea of deliberative discretion. There is a feeling that it is often up to us, in a way, whether to act on the balance of reasons when they are considerations that count for and against prospective actions in the perspective of deliberation. Failing to accord with such reasons may reflect a personal vice or a deficiency of some kind (in the dimension, say, of weakness of will), providing a potential opening for criticism and regret. But so long as our decision is supported by some sufficient reason or other, we are generally entitled to choose a course of action other than the one that is recommended by the overall balance of reasons. “Satisficing”—that is, doing something that is good enough, even if it isn’t the best option available to us in the circumstances—is sometimes a legitimate way of resolving a practical problem. When this is the case, we can parry any criticism that might be directed at us by citing the discretion we have to act in ways that are less than fully rational.

But things are otherwise with practical requirements, which are considerations that we are not similarly entitled to fail to accord with. If we are under a genuine practical requirement to do something, then there is nothing analogous to satisficing with respect to it that represents an eligible way of responding to its force. The deontic structure of practical requirements in this way contrasts with the aspirational character of the familiar reasons that count for or against many alternatives for choice.⁷ The rational force of such requirements is not something we have discretion to discount or ignore in deliberating about what to do.

We might summarize these distinctive features by saying that practical requirements enter the deliberative field in the guise of presumptive constraints on the agent’s behavior. They function in this way, insofar as their deliberative role is to determine certain options for action to be either on or off the table from the start, fixing assumptions within and around which the rest of our planning agency will operate, as it proceeds. The natural way to register such constraints in practical reasoning would be through the formation of future-directed intentions to act in accordance with them. It is a now familiar point from the literature on such intentions that they structure deliberation in a distinctive way.⁸ The fact that one intends to do X in the future is not merely one consideration among others that are to be taken into

account in ongoing reflection about action, but resolves for the agent the practical question about whether or not to do X. Further deliberation will then take place against the background of the assumption that X will be done, as the agent deliberates with an eye to resolving the other questions that are left open by that assumption.⁹

Now, sometimes we form future-directed intentions of this kind after reflection in which ordinary reasons for and against the options open to us are weighed against each other. Thus, one might decide to vacation in Venice this spring, after considering the attractions and disadvantages of doing so in comparison with the alternative of spending the same period in Reggio Calabria or Alto Adige. The point, however, is that if I am aware of a normative consideration as a practical requirement, that will itself make it reasonable for me to respond to it by forming a future-directed intention to comply, without my needing to weigh the consideration in the balance against reasons on the other side. If I have promised you I would accompany you to Venice in the spring, then I already have a basis for intending so to act, independently of how its touristic attractions compare with those of other possible destinations on the Italian peninsula at the same time of year. This is how the deliberative role of practical requirements, as constraints on agency, gets operationalized in practice.

In playing this distinctive deliberative role, however, practical requirements, it is important to note, do not impose absolute constraints on the set of options about which the agent reflects. I said above that they enter the deliberative field in the guise of *presumptive* constraints, in recognition of two different kinds of unusual circumstance that can arise. First, requirements are themselves often defeasible rather than absolute. As was noted in chapter 1, the commitment that is undertaken when I promise to do X is not a commitment to do X come what may; emergencies or other special circumstances might arise that could not have been anticipated at the time when the promise was originally made, for instance, and they can have the effect that the agent is no longer required to do X (though residual obligations, such as a duty to compensate for losses, may still obtain). Second, the original requirement, though it remains in force, might conceivably conflict with a second obligation, leaving agents in a tragic situation in which there is no way forward that respects both of the operative constraints. In this situation, options for action that were initially off the table become alternatives that agents now need to bring within the compass of practical reflection. It is not that the two requirements are to be weighed against each other in this special situation, but rather

that agents now need to face up to the fact that, no matter what they do, they will end up acting in a way that is strictly forbidden. For these reasons, we should understand the deliberative role of practical requirements to be that of placing presumptive rather than unconditional constraints on agency.

But this is a role that future-directed intentions are well suited to operationalize, for it is a point familiar from the literature on such intentions that they do not impose absolute and inflexible limits on practical deliberation.¹⁰ We form intentions from a perspective of limited information about how exactly the circumstances of our agency might change over time, but also with an implicit awareness of the ways in which those circumstantial changes might turn out to be normatively significant. This is true both when intentions are based on the recognition of a practical requirement and when they result from weighing pro and con reasons against each other. (Having decided to go to Venice in the summer rather than Alto Adige on account of my interest in its touristic attractions, I should of course reconsider if the city is visited by a natural or public health calamity in the interim.) Intentions accordingly structure the deliberations of rational agents in the manner of defeasible constraints, resolving practical questions in a way that can be revisited as new information comes in that is normatively significant. In the case of practical requirements, this might be information that shows a given requirement no longer strictly to obtain, or that brings to light a conflict with a second requirement.

The preceding considerations suggest to me that the deontic character of practical requirements represents a *sui generis* normative relation. The presumptively constraining significance that they exhibit within deliberation is distinct from, and not reducible to, the aspirational form of normativity that is at issue with ordinary recommending reasons, of the kind that count for and against candidates for action.¹¹ A normative consideration, such as a promissory commitment, can require or demand that I do something like travel to Venice, and this relation differs from the one that obtains when there are considerations, such as touristic attraction, that merely count in favor of the same action.

But this idea raises a neglected philosophical question. When we encounter a putative set of obligations or practical requirements, we should be able to make sense of them not merely as considerations that possess normative force of some kind or other, but as considerations that have the deliberative features characteristic of obligations. That is, they should not merely strike us as considerations that it would be reasonable to enter into a notional ledger of factors that speak for or against a course of action that it might be open to us to pursue.

Rather, they should be intelligible as considerations that we register from the start as presumptive constraints on our agency, constraints with which we ordinarily must comply. The philosophical challenge is to develop an account of obligations that explains their suitability to structure deliberation in the way of presumptive constraints and to support practical conclusions that are expressed using the deontic “must.” There are several familiar strategies for meeting this challenge in the philosophical tradition; let us consider, very briefly, some of the most prominent of them.

One approach, which we might call the dominance model, emphasizes the systematic importance and weight of the normative considerations that ground practical requirements. On this approach, there are ultimately just reasons for action of various strengths. What sets practical requirements apart within the larger normative domain is the fact that they can be traced to reasons that are weighty across a wide range of deliberative contexts, so that they nearly always dominate the considerations on the other side with which they might compete. It is the systematic importance and weight of their normative grounds that make it appropriate to express practical requirements with the deontic “must,” and that explain and justify our tendency to treat them as presumptive constraints.¹² If a normative reason is such that it robustly dominates potential competitors across a wide range of contexts, then we don’t really need to enter it into the deliberative ledger in order to figure out what to do, but can assume at least defeasibly that it will be dispositive. It will thus make sense to treat it as a consideration that must be complied with, one that enters deliberation in the guise of a presumptive requirement.

A different strategy appeals to the notion of identity. Some normative considerations can be traced to our self-conceptions, our sense of who we are and what is fundamentally important to us in life. In particular, there are reasons for action that are distinctively connected to threats to our identities, things that we must do if we are to hang on to our identities and to preserve them, going forward, in the face of ongoing challenges. The idea, then, is that the special nature and force of practical requirements can be traced to the special connection of the considerations that ground them to features of our practical identities.¹³ We *must* do something just in case our doing it is required if we are to fend off an existential threat to who we are and to prevent the dissolution or destruction of our selves. And this same feature explains and justifies the structural role of practical requirements within deliberation, as considerations that enter the deliberative field in the guise of presumptive constraints. It

makes sense to register normative considerations as constraints of this kind if they are grounded in the most basic features of our self-conceptions.

A third approach is voluntarism, which appeals to social relations of authority, tracing practical requirements to the commands of a suitably constituted legislator. Anscombe's "law" conception of ethics, referred to above, is an example of this strategy; it interprets practical requirements as imposed on human subjects by a benevolent and omnipotent divinity. The social relation that grounds obligation, on this account of it, is the relation that individual human subjects stand in to God. But other versions of the general approach are also possible, tracing practical requirements to the commands of human authorities, such as duly constituted legislative assemblies or to individuals to whom one is subordinate within a legitimate professional or family structure. This approach, too, identifies a feature of practical requirements that promises to render intelligible their distinctive deontic character and role. If we have been commanded to do something by a legitimately constituted authority, then it makes sense to treat it as a presumptive constraint on our agency that we will act in accordance with the authority's commands. Authoritative commands are considerations that are paradigmatically expressed using the practical "must," and they naturally structure our deliberations accordingly.

These three philosophical approaches, which I have of course sketched only in the broadest of strokes, strike me as among the most promising avenues for understanding the notion of a deliberative requirement. It is significant, for instance, that we can understand in terms of these models the common assumption that considerations of structural rationality represent requirements on practical thought. Thus, we generally take ourselves to be constrained to avoid contradictions and to take the necessary means to the ends we intend to achieve. But the constraints at issue here can be made sense of by assimilating them to one or another of the three models I have outlined. Those who believe that there are independent requirements of structural rationality, for instance, tend to treat such requirements as constitutive conditions of thought and action, which is to apply to them a version of the identity approach sketched above.¹⁴ Other philosophers reject the idea that there are independent structural requirements of this kind, arguing that the cases in which such requirements appear to obtain can better be understood as ones in which the apparent balance of reasons decisively favors a given course of action over the alternatives; this assimilates the cases, in effect, to the dominance model.¹⁵

The three models, it should further be noted, are potentially complementary. We needn't suppose that there is only one kind of factor that can potentially make sense of practical requirements, but should be open to a kind of pluralism on this issue, which allows that practical requirements might ultimately be traced to a diversity of normative considerations. Having said that, however, I now wish to outline some challenges that each of the models faces, if only to identify issues that will need to be addressed before we can arrive at a satisfactory understanding of the nature and sources of moral requirements in particular.

The dominance model, which explains practical requirements in terms of the robust weight of certain reasons for action across a wide range of deliberative contexts, appeals to a consideration that is essentially a matter of degree. Practical requirements are, in effect, systematically weighty reasons. But weight is scalar, and so it is hard to see in this account anything that would set practical requirements apart in principle from other kinds of normative consideration. The dominance model in this way seems ill-designed to make sense of the qualitative distinction that we register in deliberation between deontic considerations and the aspirational reasons that are standardly weighed against each other.

In response, it might be noted that *dominance* is not a scalar notion, even if it is constructed out of scalar materials. But it is unclear whether we can convert scalar differences into a qualitative distinction by appeal to this notion alone. Is there, for instance, a non-arbitrary way to set the threshold of greater weight by which a reason must dominate its routine competitors before it can be considered an obligation? Furthermore, this approach seems to rule out a priori something that it is important that a theory of obligation should at least leave open, namely the possibility of conflicts of obligation. In a situation that involves this kind of conflict, a single agent stands under distinct practical requirements that pull in different directions. But if practical requirements are by definition considerations that dominate the reasons with which they might compete, there cannot be a situation of this kind; the element of normative conflict undermines the very thing that accounts for the special force of obligations in the first place.

A different problem confronts the identity-based view. On this account, practical requirements are connected to the agent's identity, and one stands under an obligation to do something just in case doing it is necessary to fend off threats to some aspect or other of who one is. A strategy that takes this form, however, will have difficulty making sense of the possibility of flouted

duty, by which I mean cases in which we recognize something as an obligation, but fail to live up to it. On the identity-based approach, to acknowledge something as an obligation is to situate it in relation to a threat to one's conception of oneself, understanding that one will cease to be who one is if one does not comply with its demands. This threat, moreover, is often understood to have an existential dimension, connecting practical requirements to the core elements in one's self-conception that give one reason to go on with life in the first place. This is what is supposed to explain the appearance of peremptory necessity that sets apart those normative considerations that have the force of practical requirements. But if this is the general shape of the account, it seems to leave little space for flouted duty. If I do what I recognize will lead to the dissolution of a part of my identity, then the thing in me that is threatened cannot have the significance for my self-conception that it would have to have to ground an obligation.¹⁶ The very fact that I am willing to countenance its loss shows that it isn't among the things that contribute importantly to my sense of what my life is basically about. The only exception to this claim will be in situations in which I flout one obligation in order to comply with a second, where the second obligation is grounded in an aspect of my identity that has even greater significance for my conception of myself. But not all cases of flouted duty take this form.¹⁷

The voluntarist model seems an improvement on the other two approaches in certain respects. It traces practical requirements to a consideration that sets them apart from ordinary aspirational reasons in kind and not just in degree, and that is therefore suited to make sense of their distinctive deontic character. The voluntarist approach also allows for the possibility of conflicts in obligation (insofar as there can be commands of legitimate authorities that it is not possible for a single agent to follow), and for cases in which an agent fails to live up to an obligation whose nature and force are acknowledged. Perhaps for these reasons, it is one of the most salient paradigms for understanding the phenomenon of a practical requirement in the philosophical tradition.

But voluntarism has highly particular presuppositions that limit its applicability as a general model of obligation. For one thing, there needs to be a social relationship in place of legitimate authority, whereby one party is entitled to issue directives that govern the activities of another. The notion of authority is itself complex, and there are a variety of approaches to making sense of it that have attracted philosophical attention over the years. But however they are understood, relations of authority will obtain only under severely restricted circumstances (such as those that link the members of families and

other associations, or that imbue legislative assemblies with special normative insight or democratic legitimacy). Second, it must also be the case that the legitimate authority has actually issued a suitably public law or command in order for a practical requirement to have come into existence. There are certainly cases in which these conditions are satisfied, and they provide important examples of the phenomenon of a practical requirement. But many of the considerations that strike us in deliberation as familiar practical requirements do not derive from the public commands of a legitimate authority, but appear to obtain independently of the circumstances that are presupposed by the voluntarist model of obligation. Among these are the obligations of morality, to which I now turn.

2.2. Moral Obligation: The Specific Challenge

An account of the normative significance of morality faces some particular challenges, beyond those that are endemic to the general project of making sense of practical requirements. It must, first of all, explain why there are specifically moral considerations that have this distinctive kind of normative significance. Morality is not just a source of reasons for action, in the sense of things that count for and against prospective actions in the perspective of deliberation, but defines obligations, which enter the deliberative field in the guise of presumptive constraints on agency.

There are recent philosophical controversies about what exactly it is that traditional moral theories should be taken to be offering an account of. Do they purport to offer an analysis of the concept of the morally right or wrong, or of the property of being right or wrong in the moral sense? Or are such theories to be understood differently, for example as attempts to identify a general property that makes actions right or wrong in some other, not-yet-specified sense? Engaging with this issue, both Derek Parfit and T. M. Scanlon have recently proposed that there is an indefinable sense of “right” and “wrong,” in which to be right or wrong just is to be something that must or must not be done.¹⁸ Conclusions about what is right or wrong in this generic sense register the presence of what I have been calling practical requirements, construed as considerations that function deliberatively as presumptive constraints on agency.

I agree that moral theories are not plausibly understood, in the first instance, as accounts of right and wrong in this generic sense. They are neither analyses of these concepts, nor theories of the property that they consist in,

for the simple reason that morality is not the only source of practical requirements on the will. As we saw in the preceding section, there can be things that one must do—that it would right for one to do, in the generic sense here at issue—for identity-based reasons, and these might not have anything to do with interpersonal morality in the conventional sense. I might rightly feel I have to do something because I am someone's parent or sibling, or because acting in this way is determined by some personal project with which I am existentially identified.

But if ethical theories are not offering accounts of this generic notion of rightness or wrongness, what is their quarry? Parfit and Scanlon propose that they be understood as attempts to specify higher-level properties that *make* actions right or wrong in the generic sense.¹⁹ This seems to me a helpful way to think about at least one part of the project of moral philosophy, and I shall take it on board in what follows. But an important qualification is in order. To say that moral theories are attempts to describe high-level right-making properties (in the sense of properties that make actions “to-be-done” or “not-to-be-done”) is not to say that they are not also putting forward interpretations of rightness or wrongness in a more specifically moral sense. Thus, utilitarians, contractualists, and Kantians clearly seem to be offering competing theories of *something*, and we can usefully think of this target of analysis as a specifically moral sense of right and wrong.

Of course, proponents of these different theories propose very different conceptions of morality, organized around completely different higher-level properties (such as maximizing the good, universalizability, etc.). This can make it difficult to see them as alternative accounts of a common concept of specifically moral rightness and wrongness (as opposed to competing proposals about high-level properties that make actions right or wrong in the generic sense of to-be-done or not-to-be-done).²⁰ But I am not convinced that these are the only alternatives. A different framework for thinking about the landscape of issues here might start by identifying an abstract concept of morality, marking out a domain of things that are right and wrong in a specifically moral sense. This domain collects reasons and values that have to do, in some way or other, with the social significance of our actions for the individuals who stand to be affected by them; it is the domain that I characterized in chapter 1 as interpersonal morality.

Our sense that there is a coherent domain of interpersonal reasons and values reflects our acknowledgment that there are different kinds of objections to what a person might propose to do, and that not all of these are of a moral

nature. There is something morally problematic about taking advantage of other persons' vulnerability, or deliberately deceiving them about a matter that is important to them; by contrast, actions that can be criticized for reasons of style or efficiency or personal conviction need not be morally problematic in the same way.²¹ The conventional concept of the moral right, in this abstract sense, might itself be historically conditioned; indeed, it seems to me plausible to suppose that the concept most of us operate with today is a distinctively modern one, incorporating elements that might have been lacking in other places and times. But for better or worse, it is our concept of the moral right, and it includes abstract features that constrain theorizing about morality in the narrower sense at issue. Moral theories can then be understood to put forward substantive conceptions of the moral right and wrong in the sense that is defined by this modern, interpersonal concept of morality.²² Such theories are in competition with each other, insofar as they are attempts to do justice to constraints to which any account of interpersonal morality is generally understood to be answerable. But they differ from one another, insofar as they fill in the general concept in very different ways, identifying different properties with moral rightness and wrongness, and offering us different ways of understanding what it is to correctly think of something as morally right or wrong.²³

Applying this framework to the immediate task of the present chapter, it is one of the abstract constraints on a substantive conception of interpersonal morality that it identify a high-level property that makes actions right or wrong, in the generic sense of to-be-done or not-to-be-done. This is just to say that theories of morality must make sense of the idea that there is a morally significant high-level property that correctly figures in deliberation as a practical requirement or an obligation. Meeting this desideratum, I maintain, is a distinctive and underappreciated challenge. It requires us to make sense of the properties specified by a conception of moral rightness as properties that are not merely normatively relevant to deliberations about action, but relevant in the distinctive way of practical requirements: as considerations that properly register in deliberation in the guise of presumptive constraints on agency. Of course, it is possible that this theoretical challenge is one that cannot be met. Perhaps there is no interesting high-level property that is held in common by all of the actions that we intuitively think of as morally right and wrong. Or perhaps, though there is such a property, we cannot make sense of it as one that is significant in the way of practical requirements. These possibilities point toward a skeptical or debunking account of interpersonal morality, one that denies that morality represents a unified domain of obligations. My point,

however, is that a nondebunking account of moral rightness will offer a conception of it that illuminates its deontic structure, its standing as a source of practical requirements of the kind discussed in section 2.1.

A second feature of morality that I should like to emphasize is one that is characteristic of the modern concept of interpersonal morality in particular. I noted above that the concept of morality in the narrow sense collects reasons and values that connect to the social dimension of agency. But there is a specific understanding of this social aspect that is salient in modern moral thought. There are different ways to express the point, but one of them would be to say that morality defines a cosmopolitan normative structure. The moral community, according to this conception of interpersonal morality, is a maximally inclusive group of individuals whose interests are taken to matter equally. Those who are accorded equal standing by morality include, at a minimum, all human beings, though different views can be taken as to whether and in what ways the moral community extends even further, beyond the members of this class. One such further extension would be to those beings, if there are any, who share with normally developed humans the capacities for practical reason and reflective self-determination; another would be to creatures with interests that are grounded in their having a coherent point of view and a capacity for pleasure and suffering.

The general idea is that, from the moral point of view, all members of this extensive class, however precisely its limits are defined, are to be taken into account, as beings who are neither more nor less important than the other individuals in the class. This postulate of equal standing, as I referred to it in section 1.2 above, is perhaps the defining moral insight of modernity, and it has a bearing on the modern concept of morality that different theories are attempting to give an account of. Thus, the normative significance of morality for individuals is connected directly to the fact that we are members of an inclusive community of equals. On the characteristically modern way of understanding it, interpersonal morality might be thought of as the solution to the problem of how each of us is to negotiate a social world of this cosmopolitan and inclusive kind, and its verdicts matter to us, as agents, in virtue of its playing this important role.

Putting this idea together with the first, we could say that morality is the set of deontic constraints on conduct that derive from the fact that we inhabit an extensive notional community together with other beings who are “equally real” (in Thomas Nagel’s striking formulation), and whose interests are no less significant than ours.²⁴ A philosophical account of interpersonal morality must

shed light on these twin aspects of the moral, and the connection between them. Its substantive conception of moral rightness must show that the property of being morally right or wrong is not merely a presumptive constraint on the rational will, but one that reflects our common membership in an extensive community of moral equals. It should thus fall out of a substantive conception of interpersonal morality that in treating moral rightness or wrongness as practical requirements, we are acknowledging our equal standing within such a moral community.

The third feature I will mention is somewhat more controversial, and perhaps less central to the modern concept of morality. But it is very deeply embedded in our first-order moral thinking, and it is therefore a desideratum that it is reasonable for moral theories to attempt to accommodate. I have in mind here the idea that moral requirements, at least in the central cases, have an agent-relative character. That is, they do not merely define impersonal or agent-neutral values that we take it to be important that all people should promote in one way or another. Rather, they are requirements that assign to different agents different moral aims.²⁵ Thus, the moral significance of promissory undertakings reflects itself in the fact that we are under a presumptive obligation to keep the promises that we ourselves have entered into. It is a misunderstanding of the relevant moral desideratum to think that the fundamental value that is at stake is the impersonal value of promissory fulfillment, and that we respond correctly to it by doing everything in our power to see to it that it is maximally realized in the world. We are morally obligated to keep our own promises, even if we happen to inhabit the philosopher's distressing scenario in which our doing so will lead several other agents to break promises they would otherwise have kept.

These three features, I shall assume in what follows, are reasonable constraints that a substantive account of the moral right should aspire to meet. But it is a serious challenge to devise a theoretical account of interpersonal morality that does justice to these three constraints together. To illustrate the challenge, it may help to consider very briefly how some familiar accounts of morality would approach it.²⁶ I shall suggest that we can see these accounts as exploiting the different models of a practical requirement that were surveyed in the previous section; but also that those models cannot easily be married to the other specific features of moral obligation that I have just enumerated. The result is one that Anscombe anticipated, namely that moral obligation continues to seem puzzling if we think of it along the lines of the most salient modern moral theories.

Start with classical utilitarianism, which holds that those actions are morally right that maximize the net balance of pleasure over pain (producing at least as much hedonic utility as any of the alternatives that it was open to the agent to perform). As I noted in section 1.2, an approach of this kind provides an appealing account of the cosmopolitan aspect of morality. Deliberating from the Sidgwickian point of view of the universe, we are to take equally into account the interests of all of the sentient beings potentially affected by our actions, operating with austere impartiality as between those individuals. Moral requirements thus derive in a straightforward way from our acknowledgment of our equal standing with others as members of a common moral community. But if utilitarianism does well with this feature of morality, it does less well with the other two I have identified. Whatever else it might be, the principle of utility is the antithesis of an agent-relative principle; it enjoins us to promote a conception of value that is fundamentally agent-neutral in character. As long as the outcome that we bring about is impersonally optimific, it is a matter of indifference to utilitarian morality what the character is of the action that we perform. There is no room on this picture for agents to take a special concern for their own veracity or loyalty or fidelity to agreements. There is something perhaps bracing about this way of thinking of morality, but it is undeniably at odds with naïve ways of understanding what interpersonal morality asks of us.²⁷

Nor does utilitarianism yield a plausible story about the deontic character of moral considerations. An action is morally right, on this approach, if it would produce at least as much total pleasure or happiness as the alternatives that it was open to the agent to choose. But why should an action's being right in this sense be a consideration that enters the deliberative field as a presumptive constraint on agency? As we saw above, there are different models for thinking about practical requirements, but utilitarian rightness doesn't seem to fit any of them. For instance, according to the dominance model, an action might be one that we are intuitively required to perform if there are decisive reasons that speak in favor of doing it, across a wide range of deliberative contexts. When this condition is satisfied, it makes sense to treat the action as one that is defeasibly fixed for purposes of future deliberation and planning, since we can be confident that the balance of reasons will continue decisively to favor the action as we progress through new situations.

But utilitarian rightness doesn't seem to function this way, for two reasons. First, to say that a course of action would produce the best consequences is not to say that it is decisively supported by the balance of reasons, even if we

grant the hedonistic assumption that our reasons are exclusively tied to the production of pleasure and pain; the action that is best in these terms, after all, might merely produce a slightly better balance of pleasure over pain than one of the alternatives open to the agent. The utilitarian conception of rightness, considered on its own, thus leaves it open whether the course of action it recommends is decisively supported by the balance of reasons in any given case. Second, the action that is right in one situation, in utilitarian terms, might turn out in the very next situation one encounters to be less productive of utility than an alternative that is available in that situation. It all depends on the consequences, as the utilitarian would say. But then it wouldn't really make sense to treat the right action in this sense as something that one is committed to doing, going forward.²⁸

The utilitarian might respond that this last problem stems from a fixation on the wrong level of action description. The thing that one can be confident will always be supported by the balance of reasons is not an action of some independent type that happens, in a given situation, to be the right thing to do (say, keeping one's promise or assisting a stranger in need). It is, rather, doing what would produce the best consequences, under that very description. But this does not represent a consideration that it would be intelligible to treat as a presumptive constraint on agency of the kind that we have seen practical requirements to represent. In any situation we might find ourselves in, if we are to act in accordance with the utilitarian construction of moral rightness, we will have to give some thought to the particular consequences of the options that are open to us, under the particular conditions that then obtain. Nothing is antecedently off the table, so to speak, not even in the provisional way that characterizes the deliberative role of obligations. Utilitarian rightness thus does not seem well suited to define a self-standing class of practical requirements.

We saw earlier that the dominance conception faces the difficulty of trying to extract a qualitative distinction, between ordinary reasons for action and practical requirements, from something that is essentially a matter of degree. A similar difficulty faces the utilitarian conception of the right. That conception, couched as it is in scalar terms, defines an ideal of attainment that agents could aspire to live up to, as far as it is possible for them to do so, much as the devotee of a religious doctrine might aspire to approximate to a maximally exigent standard of personal purity or devotion. Scalar ideals of this kind represent intelligible objects of personal ambition. But it is not clear that it really makes sense to treat compliance with such a scalar ideal as a presumptive con-

straint on one's deliberations and activities as an agent.²⁹ Moral rightness, as utilitarianism defines it, seems better thought of as an aspirational than as a deontic notion. It is commonly observed that the utilitarian conception of the right doesn't leave any room for the supererogatory; perhaps this is because it equally doesn't deliver a credible notion of the obligatory.³⁰

Of course, refinements are possible, perhaps drawing on different intuitive models of an obligation. We saw in section 2.1, for instance, that practical requirements are often thought of in voluntarist terms, as connected to the legislative acts of authorities about the things that it is open to us to do. In this spirit, we might, following Mill, characterize moral obligations as those moral standards that it would be optimific to sanction people for flouting, interpreting the attachment of sanctions to standards as a kind of inter- or intrapersonal legislation.³¹ This would arguably represent a reconstruction, in utilitarian terms, of something that is recognizable as a conception of moral obligation. It makes sense for agents to commit themselves presumptively to ensuring their compliance with standards to which internal or external sanctions are in this way attached. But the resulting account would abandon the aspiration to explain why moral considerations have the intrinsically deontic character of demands on the rational will. Thus, from the fact that an action is morally wrong, it will not follow on this approach that the agent is under an obligation not to do it. Whether that is the case or not will depend on ancillary considerations, about the consequences of a general scheme of moral sanctions that attaches penalties to people when they perform actions of the kind that is in question. The utilitarian account of what makes actions right or wrong thus does not succeed in delivering an explanation of the idea that acts that are morally wrong are ones that we have an obligation not to perform.³²

To continue in this cartoonish vein for a moment, consider next a more robustly voluntarist position, according to which what makes actions right or wrong in the first place is the fact that they are required or prohibited by the commands of an authoritative lawgiver of some kind. On this venerable approach, there are a variety of desirable ways for people to act and to interact with each other, but it becomes right or obligatory to act in those ways only through the laying down of a corresponding requirement by an appropriately constituted authority. As noted earlier, Anscombe's critique of modern moral philosophy seems to presuppose this conception of obligation; her charge is roughly that the modern philosophical approach to morality posits universal obligations that lack the foundation in the commands of a divine authority that their intelligibility as obligations requires.³³ But there are also secular

variants of the strategy, which trace requirements to the commands of a human legislator or authority.

An approach that explicates the morally right in terms of such authoritative legislative acts has some undeniably attractive features. As I observed in section 2.1, the voluntarist model, at least when its presuppositions are in place, provides an intelligible paradigm of a practical requirement, so a voluntarist theory of rightness promises to make sense of the deontic structure of the moral. The approach also seems capable of accommodating the agent-relative character of familiar moral requirements. If God, or the secular authorities, command individuals to keep the promises they have entered into, then one would be flouting the resulting requirement rather than complying with it if one were to bring about a net increase in the incidence of promise-keeping by breaking a promise one had made oneself. That is just not the thing that the relevant authority commanded one to do.

But the cosmopolitan aspect of modern morality is difficult to make sense of in voluntarist terms. The challenge here was to explain how moral rightness has its source in our common membership in an extensive community of equals. But if voluntarism takes the most common form of a divine command theory of the moral right, this desideratum seems to go by the board. The requirements that God lays down on us may enjoin us to treat other individuals with consideration and respect, as moral equals in some sense or other (perhaps as persons who are equally subject to the divine will). But this is a matter of the content of what is commanded, not of its normative status as an obligation. What makes it the case that we are practically required to comply with an injunction that has this content is solely the fact that a benevolent deity has commanded us so to act. God might equally have commanded us to comply with laws whose content has nothing to do with the equality of individual human subjects (such as dietary restrictions that prohibit the consumption of certain flora and fauna), or that even deny such equality. Or God might have imposed on the members of one tribal group requirements that are not similarly imposed on the members of other tribal groups, singling out the first community as a kind of chosen people. The divine command approach, construed as a way of understanding moral requirements, puts into the foreground the relationship between God and the individuals who are subject to God's laws, rather than the relationships that members of the broader community of moral agents stand in to each other.

It is perhaps not surprising that divine command theories have trouble accommodating what I have called the cosmopolitan aspect of morality, for it

represents a distinctively modern desideratum, one that helps to define the Enlightenment project of understanding interpersonal morality as a secular phenomenon. But the problem remains intractable when we turn to voluntarist views that invoke a human rather than a religious authority. One such view holds that moral rightness is to be understood by reference to the expectations that are actually imposed on us by the members of the communities in which we happen to live. It is morally right, and hence obligatory, for me to do something, according to this approach, just in case other people in fact demand that I so act, where demands in the relevant sense are associated with informal social sanctions (including angry disapprobation and the like).³⁴ This social command theory, as we might think of it, borrows some of the advantages of the divine command theory for making sense of the deontic structure of the morally right. But the way that the voluntarist apparatus gets transposed into a secular key by this approach seems too crude to yield a plausible account of moral obligation. The expectations that are actually backed up by attitudinal sanctions in human communities are often misguided, reflecting prejudice, ignorance, and superstition as much as moral insight; indeed, they can operate at cross purposes to ideals of equality, insofar as there are communities in which norms of racial or gender supremacy and oppression are enforced through the local economy of disesteem and social sanction. A different way to put the point might be to say that the members of our actual communities do not have the kind of default authority that would make the commands they might direct toward us automatically binding.³⁵

This difficulty might be avoided by tracing obligations, not to the actual commands of our fellow citizens, but to the fact that we all have “*de jure* authority” to address moral demands to other agents in our guise as representative members of the moral community.³⁶ This represents a shift from actual to hypothetical commands, together with the invocation of an idealized standpoint that each of us can in principle occupy for issuing commands on behalf of everyone. These modifications in the social command theory introduce some critical distance between the potentially misguided expectations of many actual communities and the genuine requirements of morality, while also giving expression to the ideal of an inclusive community of equals. But they also undermine the presuppositions that make voluntarism viable as a way of modeling practical requirements. For one thing, the notion of “authority” that results when voluntarism is revised in this way seems entirely vestigial. We are left with the idea that actions are obligatory when there are reasons for any person to address the corresponding demands to the agent;³⁷ but then

it is these reasons, rather than any authority that might be invested in us, that are doing the real explanatory work. For another thing, this modified approach abandons an assumption that seems central to the paradigm cases of authority-based requirements, namely that the legislator should exercise its authority by issuing an actual command. Parliament may possess the *de jure* authority to modify the tax code in the direction of imposing higher rates on the affluent, and there might even be good reason for it to do so; but its subjects are not required to comply with such notional rates if the corresponding legislation has not actually been enacted. Merely hypothetical laws are not in fact binding on anyone.³⁸

Some of these difficulties could be avoided by internalizing the relation between authority and subject, treating moral and other practical requirements as laws that agents legislate for themselves. This recognizably Kantian approach transforms voluntarism into an ethics of autonomy.³⁹ The moral law is binding on each of us insofar as we impose it on ourselves, through commitments that are built into the structure of free agency. It is striking that Kant in this way adapted the voluntarist model to make sense of morality as a source of nonnegotiable demands on the rational will; his interest in doing this shows the power of the voluntarist paradigm when it comes to understanding what it is for something to be a practical requirement.

It is equally striking, however, that a Kantian ethics of autonomy takes on board some very ambitious philosophical assumptions, ones that have not been defended or developed persuasively in over two centuries of work that is broadly inspired by his ethical vision. There is a need, for one thing, to show that the voluntarist model can coherently be applied in thinking about the relation that agents stand in to themselves. (What is the nature of the authority that is in play here? How can any practical law be binding on us if it is always open to us, as legislators, to modify it if we no longer wish to comply?) And there is a need, second, to demonstrate that free agents really are committed to imposing on themselves the moral law.⁴⁰ Only this will ensure that moral obligations apply universally, to all agents who are members of our notional cosmopolitan community of equals. I am not optimistic that philosophers working in this tradition will be able to vindicate these important assumptions.

Consider, finally, a perfectionist approach to basic moral requirements, one that derives them from considerations about what Philippa Foot has called “natural goodness.”⁴¹ An action is right, on this approach, if its performance is required by traits that people need in order to be good human beings, and

wrong if it is incompatible with such traits. Thus, we start by thinking about the challenges that human communities typically face under the conditions that they normally inhabit, and identify a set of characteristics or virtues that members of such communities need to instantiate if they are all to do well under the conditions specified. Individuals can then be said to be good as humans if they have the traits that people generally have to have if they are to flourish under the circumstances that they typically encounter. And actions will be right or wrong, we could go on to say, if doing them is either determined or ruled out by the traits that make people good humans in this sense. Facts about natural goodness thus provide the template for understanding normative claims about individual living creatures, including claims about what it is right or wrong for those individuals to do.

This interesting approach raises a host of foundational questions. Abstracting from the details, however, we can see that part of its appeal lies in its promise to shed light on at least two of the elements in our thinking about morality that I have highlighted. The deontic character of moral considerations might be explicated in terms of the identity-based model of a practical requirement sketched quickly in section 2.1. Living things are members of a life form, whether they like it or not, and this determines standards of conduct that are not optional for them, but that flow from their essential identity. The wolf that hangs back from the pack, rather than contributing with vigor to running down the prey they are hunting together, is defective as a wolf, even if the result turns out to be advantageous for it on this occasion. Similarly, the person who acts wrongly will be failing relative to standards of attainment that are not optional for human beings, insofar as those standards are associated with the virtues that humans generally need in order to flourish under their normal conditions of life. Individual agents have to comply with the standards of virtue, we might say, on account of who they are, and those standards therefore represent practical requirements for them.⁴²

The approach also renders intelligible the agent-relative character of basic moral requirements, as we ordinarily understand them. What I as an individual am required to do is determined fundamentally by the virtuous traits that I have to exhibit to count as a good human being myself. So the important thing is that I should be just and generous and reliable, not that my actions should contribute causally to the maximal instantiation of these virtues in the population at large. I must keep my own promise, even if my breaking it would lead several others in the local community I inhabit to keep promises they would otherwise have broken.

Like voluntarist accounts, however, perfectionism does less well with the cosmopolitan aspect of morality that I have maintained is central to the modern conception of it. The idea, again, was that the requirements of morality have their distinctive source in the fact that we inhabit a common world together with others who are equally real, constituting a maximally inclusive community of equals. This idea seems to drop out of the picture, however, on the perfectionist approach. Moral requirements, to the extent there are such things, trace their origin not to the direct significance of other peoples' interests for us, but to the value of individual perfection. The basic question is not how we are to negotiate life as members of a community in which other peoples' interests are no less significant than our own, but how we are to realize an ideal of human attainment in our own case.

True, the characteristics we need to exhibit in order to be good as individual human beings will include such virtues as justice and generosity, which have to do in part with how we relate to and interact with other people. Furthermore, people who have these characteristics will perform just and generous acts for their own sakes, not as a result of reflection on what it is to be good as a member of the life form to which they naturally belong. One consequence of this is that the character of the practical requirements defined by this approach is not fully salient in the deliberations of those to whom the requirements apply. We may be subject to a requirement to act in ways consistent with the virtues; but the basis of this requirement, in considerations about what it is to be good as a member of the life form to which we naturally belong, is one that will ordinarily be transparent to us in deliberation. A further consequence is that, at the level at which normative requirements are explained, the interests of other people enter as occasions for the realization of virtue, rather than direct sources of requirements on the virtuous agent. Their significance for questions about what the agent is required to do is thus mediated by the demands of human goodness, which is taken to be the ultimate source of normative requirements.

This shows itself in the fact, emphasized by Foot and other proponents of a broadly Aristotelian approach, that the virtues one has to have in order to achieve natural human goodness include traits that have nothing to do with the needs and interests of other persons, such as a hopeful outlook and an ability to accept good things in one's own life when they come one's way.⁴³ The sense in which one is under a requirement to adopt these attitudes, on the approach in question, is just the same as the sense in which one is required to

treat people justly or with compassion; which is to say that the requirements at issue are not derived immediately from facts about our common membership in an inclusive community of equals.⁴⁴

2.3. A Relational Approach to Moral Obligation

The remarks of the preceding section do not of course amount to a comprehensive critique of the traditional moral theories that were under discussion. My intention was instead to illustrate the specific challenges that face an account of moral obligation, and to identify the most salient weaknesses of some standard approaches to morality when it comes to meeting those challenges. For all I have said so far, however, it might be that one of the theories in question could be modified or adapted to provide an improved philosophical account of moral obligation. Alternatively, an account that seems deficient in light of one or the other of the challenges I have identified might nevertheless represent our best hope for making some sense of moral obligation. To identify the challenges that confront the project of understanding this phenomenon is not necessarily to take for granted that it will be possible to meet all of them. We might in the end conclude that the best we can do is to render intelligible some but not all of the features that we intuitively associate with the notion of moral obligation. Or we might, more radically, join Anscombe in concluding that the notion has no coherent place in a modern conception of ethics, and that we would be better off doing moral theory without it.

In the present section, however, I would like to move in a more optimistic direction, presenting in outline an approach to moral obligation that makes better sense of its central features. My point of entry will be an interesting moment in Foot's development of the natural goodness approach. In the course of sketching her version of a perfectionist theory of ethics, Foot takes up the question of our reasons to keep the promises we have made, focusing on the case mentioned in chapter 1 of Mikluko-Maklay's promise to the Malay servant. In this case, it will be recalled, the anthropologist Maklay promises the servant not to photograph him, where the servant extracted this promise from the anthropologist in the belief that his spirit would be harmed if a picture were taken of him. Foot takes it as given that the promise changes the anthropologist's normative situation, and she wonders about the nature of the reasons that it brings into existence.

Now it might have been thought that the answer to this question would be pretty straightforward for a theorist of her convictions. We need only identify some virtues—that is, traits that humans generally need in order to do well under what for them are the normal circumstances of life—that would be incompatible with breaking the promise, despite the lack of harm it would in fact cause the servant. And indeed, Foot herself mentions some such virtues in this connection, including trustworthiness and respect.⁴⁵ We could then say that a good human being, someone who exhibits such virtues as trustworthiness, will take the fact of the promise as a reason to fulfill its terms, and that this is sufficient to illuminate its normative significance.

But Foot apparently finds this answer inadequate to account for the special force of the reason that is introduced by Maklay's promissory undertaking.⁴⁶ To make sense of this aspect of the case, she suggests, we need to observe that promises exploit "a special kind of tool invented by humans for the better conduct of their lives, creating an *obligation*."⁴⁷ The idea, though underdeveloped in Foot's presentation of it, is apparently that when Maklay thinks that he should not photograph the servant, he understands the promise to have brought into existence a new and specially exigent normative fact, which we could refer to as an obligation. He has bound himself through his promise; and there are presumably further natural virtues that kick in (such as the virtue of fidelity) to explain why good human beings would not do something that goes against what they take themselves to be in this way obligated to do.

I am not sure that it is helpful to think of obligation, in the sense that is at issue here, as a "special kind of tool" that humans have invented. But Foot is certainly correct to observe that promises are generally understood to give rise to distinctive kinds of reasons, ones that are in the key of practical requirements. She also seems to me to be correct in connecting their distinctive force to the fact that they are ways of binding ourselves to do something, which bring into existence what we intuitively understand to be obligations. Indeed, these connections seem to me to point the way toward an improved solution to the general problem of understanding moral obligation. The solution will begin by taking seriously some of the distinctive features that are at work in the case of promissory duty, and treating them as paradigmatic for the more generic phenomenon of moral obligation.

Foot's basic thought seems to be that we understand ourselves to be bound or obliged when we promise to do something. The sense of obligation that is at issue in such cases, I would submit, is roughly the relational conception of normativity whose elements were sketched in chapter 1 of this book. Maklay's

thought, fully spelled out, might be that he owes it to the servant to refrain from taking photographs of the servant's person, because he made a promise to the servant that he would respect the servant's wishes in this matter. There is a normative nexus, consisting of a directed duty and a corresponding claim or entitlement on the part of the Malay servant, that has been created through the promissory act, and its presence is crucial to understanding the difference that the promise makes to their relations with each other from that point on. This shows itself, for instance, in our understanding that if Maklay were to break his promise to the servant, he would not merely have acted wrongly, or displayed a deficiency that undermines his status as a good human being; more specifically, he would have wronged the servant in particular.

A directed obligation of the sort at issue in this case plainly has at least some of the features that we were looking for in an account of the normativity of the moral domain. It is, for one thing, the kind of consideration that is plausibly understood to exhibit the deontic rather than the merely aspirational form of normative significance. We might say that obligations, in the specific sense of relational or directed duties, are among the considerations that make sense to us as practical requirements in the generic sense. An action can be "to-be-done" or "not-to-be-done," just insofar as and just because it is something that we owe it to another party to do or to refrain from doing. Obligations in the relational sense are thus paradigmatically intelligible to us as practical requirements (or "obligations" of the more generic kind); this is, in large part, what makes them normative notions in the first place.

As we saw in section 2.1 above, the difference between practical requirements in the generic sense and aspirational reasons is in part a matter of differences in regard to deliberative discretion. We have the sense that many of the reasons that count for and against prospective courses of action are considerations that we have some leeway to discount in our practical thinking about what to do, whereas a similar discretion is out of place in regard to practical requirements. But the relational aspect of directed obligations helps us to make sense of this contrast. In a case with the inherently relational structure I have described, one's reasons for doing something are constitutively connected to claims to performance on the part of another person. The values in which these reasons are based are not purely monadic; they do not exclusively concern the agent, but essentially implicate the person to whom the agent is related, whose own normative situation will be altered if the agent fails to respect the value that is at issue. We might say that they are held in common by two different parties, the agent who stands under the directed duty and the

claimholder to whom it is directed. But when these features are present, it seems natural that the agent would lack the unilateral discretion to discount the normative consideration that seems to be present with some other kinds of reason for action. The reason in question is part of a normative complex that does not belong to agents alone, and so it isn't for agents to determine how much weight it is to be given in their practical deliberation.⁴⁸ "Satisficing" with respect to it is not an eligible option.

A further dimension of practical requirements, I suggested earlier, is their function within deliberation as Raz-style exclusionary reasons. They are not considerations that belong in a notional ledger of pros and cons, as items that are to be weighed against other considerations of the same kind. Rather they enter the deliberative field as presumptive constraints on the agent's behavior, determining that some things are provisionally to-be-done, and that others are provisionally off the table. Practical requirements that structure deliberation in this way play a very different role from the ordinary normative reasons that count for and against prospective options that the agent might pursue, and it is a general challenge for a theory of obligation to make sense of this important dimension of them.

But the directed duties that are created by promissory exchanges, such as that of Maklay, are normative considerations that intelligibly function as presumptive constraints on agency of just this kind. Transactional duties, including those created by promises and other forms of agreement or exchange, represent what is often thought to be the original notion of obligation: an obliging of one agent by another, which brings into existence a normative debt that must be repaid.⁴⁹ The resulting obligation is something that we are aware of in deliberation, as a consideration with its own normative significance; having made a promise to do X, or signed a contract so to act, I will naturally think, going forward, that I now owe it to the other party to fulfill the commitment I have undertaken.

But it is characteristic of such *commitments* that they are properly understood to function as presumptive constraints within the agent's practical reflection about what to do. To undertake a commitment of this kind to someone is ordinarily to return a provisional answer to at least one of the questions that might be asked about how one is going to comport oneself, going forward. Having made a promise to do X, I will naturally think that I now owe it to the promisee to fulfill the promissory commitment. I will therefore take my doing X to be a presumptively fixed point in my ongoing planning about my activities, forming an intention to live up to the commitment that is in this way owed

to the other party. The directed obligation is thus something that intelligibly impinges on deliberation in the way of a practical requirement; it functions to rule out certain options (at least defeasibly), and to determine that others will be performed, in virtue of the connection of those options to the claims of another individual.

Transactional duties, such as those generated through promissory exchanges, thus seem to exhibit the deontic features that we have seen to be characteristic of practical requirements. In virtue of their constitutive connection to the claims of another party, they are considerations that we don't have natural discretion to discount, and they rightly function in deliberation as presumptive constraints. Whether or not they represent the original notion of an obligation, they are certainly a natural paradigm for this distinctive phenomenon, and they are widely understood as such. This is especially the case if we assume, as I have been doing, that promissory debts are morally dispositive, representing normative commitments, rather than merely *pro tanto* considerations that enter into determining what we might be obligated (in some yet unspecified sense) to do. Promissory duties might be defeasible, but when they obtain, they define moral obligations that are binding on us, precisely insofar as they specify things that are owed to another party.

But what about the other features that we have seen to be characteristic of specifically moral obligations? Can we make sense of these in terms of the relational model that is exemplified in cases such as that of Maklay's promise? Consider the issue of agent-relativity. Moral obligations, I have suggested, are familiarly understood to exhibit this distinctive kind of structure, but it is a feature of them that has often been thought to be puzzling or even paradoxical. How can a concern for a basic value lead me properly to care about the instantiation of the value in my own life, without caring equally about its instantiation in the lives of others?⁵⁰ Yet this is precisely the structure of concern that seems to be determined by the ordinary understanding of moral obligations. Here we might begin by observing that this puzzling structure appears to be latent in the examples of transactional duty that we have been considering in this section. Thus, having undertaken a promissory obligation, I now owe it to the promisee that I in particular should do the thing that was promised. The promissory transaction does not generate an impersonal concern to promote the value of promissory fidelity in the larger community of moral agents, nor can it plausibly be derived from an agent-neutral concern of that kind. Intuitively, at any rate, transactional duties seem to exhibit the characteristically agent-relative pattern of concern.

Moreover, it appears that we can explain this agent-relative character straightforwardly by appeal to the relational content that we have seen to be characteristic of these obligations. The promissory transaction creates a nexus between the two parties to it, the promisor and the promisee; the former owes it to the latter to do what was promised, and the latter has a claim against the former that the promisor should so act. The obligation that someone like Maklay has undertaken in making a promise is thus one that can be honored only by remaining within the relational compass that it in this way defines—that is, by acknowledging the Malay servant's claim, and thus living up to what he owes it specifically to the servant to do. Other agents may enter into similar transactions with each other, giving rise to directed duties and corresponding claims that link them in their own normative nexus. But seeing to it that they honor the duties that they owe to each other is not something that Maklay is responsible for, either in general or in virtue of his having pledged his word to his Malay servant.

Granted, it would formally be possible to make a promise whose content was to advance some agent-neutral goal in a particular domain of action. I could promise you that I would undertake to maximize impartial welfare in my decisions about how to expend a given percentage of my net income, going forward. I might even promise you that I would do what I could to maximize the incidence of promissory fidelity among the members of my broader community. But these agent-neutral goals would be grounded in the promissory commitment that I have undertaken specifically to you, which is a commitment that I should comport myself in the way I have promised to do, by pursuing the agent-neutral goals specified in the promise. If somebody else could make it the case that those very goals would be better realized through my own personal failure to pursue them, this would not count as a way in which I had fulfilled my promissory commitment. (The situation might be different if I induced the other person to bring about this state of affairs, or accepted an unsolicited offer from the person so to act; in these scenarios, my initial act of getting the other person to intervene would count as honoring my promise to you, even if it had the effect that I would, in future, no longer be pursuing the agent-neutral goal to which I had committed myself.⁵¹) If we understand moral obligations along the lines of the model implicit in the promissory case, then the agent-relative character of them seems to fall out fairly directly, as a fully intelligible consequence of their inherently relational structure.⁵²

Let us turn, finally, to the cosmopolitan aspect of modern morality. The idea here was that specifically moral obligations have their source in the circum-

stance that we inhabit an inclusive world of moral subjects who are equally real. We understand ourselves to be members of a maximally extensive community of individuals whose interests are no less significant than our own, and our moral obligations derive from this fundamental fact. Is this something we can make sense of on the relational approach to obligation that I have been exploring in this section?

The first thing to note about this question is that the paradigmatic examples of relational duties that have so far been in the foreground have features that are in tension with the cosmopolitan aspect of the moral. Transactional duties, such as Maklay's promise to the servant, rest on a causal interaction between the two parties that they bind. But it is plain that we do not interact causally in this way with all of the members of the maximally inclusive community of moral subjects. If we think of moral obligations in essentially transactional terms, then, we will clearly not be able to do justice to the inclusive aspect of morality, as I have so far presented it.

Suppose, however, that there are directed obligations that do not rest on specific exchanges or transactions with other individuals, but that specify what we owe to people just in virtue of the fact that they occupy a world in common with us, and are therefore liable to be affected in one way or another by the things that we might decide to do. Just as promisors owe it specifically to promisees to live up to the commitments they have entered into, so too might there be things that we owe to each of the other members of this maximally extended moral community, regardless of our antecedent relations to them. This would be a cosmopolitan conception of morality as an interpersonal domain of distinctively relational obligations. The cosmopolitan conception might be thought to result from generalizing the original, transactional model of an obligation to encompass all the members of an inclusive community of equals, linking them pairwise in a distinctive kind of normative nexus.

For reasons that should now be plain, an account of this kind would seem exceptionally well situated to illuminate the specific normative features of the moral to which I called attention in the preceding section. It would represent morality as a set of directed obligations, which we correctly register in deliberation as presumptive demands or requirements. Furthermore, these obligations would have an intelligibly agent-relative character, which is determined straightforwardly by their inherently relational structure. Finally, the resulting picture would offer an appealing account of the cosmopolitan aspect of morality; it would connect moral obligations constitutively to claims that others have against us, simply in virtue of the fact that we and they are

members of an inclusive community of individuals whose interests are equally important.

In my view, much of the attraction of the relational interpretation, as a unified interpretation of the moral domain, lies in its promise to illuminate these basic normative features of interpersonal morality, which as we have seen are difficult to make sense of on other approaches. Of course, there is much that remains to be done before we can be confident that this relational interpretation of the moral is genuinely viable. I have so far merely invited the reader to suppose that the transactional model of obligation might be generalized to define a set of duties that are owed to anyone who might stand to be affected, in one way or another, by what we do. But there are large philosophical questions that are raised by this suggestion, which will need to be tackled before we can be confident that the cosmopolitan version of relational obligation is really defensible. We will also need to explore the first-order normative implications of understanding the domain of interpersonal morality in these relational terms. These are tasks that I take up in chapters 4 through 6 of this book. For the remainder of this chapter and the one to follow, however, I shall operate on the assumption that the generalized version of transactional obligation represents at least a possible template for understanding interpersonal morality, in order to continue to develop the theoretical advantages of this approach to interpreting the moral right.

2.4. Refining the Picture

Assume, then, that the transactional model can be generalized, so that we are linked in a normative nexus to each of the individuals who might potentially be affected by the exercise of our agency. Every one of these individuals would have moral claims against us, just insofar as they have standing as equal members of an inclusive moral community. I have so far argued that an approach along these general lines would be well suited to shed light on some of the specific features of moral obligation, illuminating the distinctive normative significance that moral considerations seem to have in the first-person perspective of deliberation. But to understand the moral in these terms, we will need to take on board some controversial and hitherto neglected commitments of the approach.

The relational account represents moral considerations as obligations in the original sense, specifying duties that are owed to the various individuals who

are in a position to be affected in one way or another by what we do. According to this conception, the highest-level property of moral rightness just is the property of being owed to another party or parties, insofar as they are persons with equal standing. And to correctly identify something as morally right is to think of it as having this property, so that one conceptualizes oneself as linked to another person through a network of connected claims and directed duties.⁵³ Considerations of this relational kind enter the deliberative field in the guise of presumptive constraints on agency, in ways that align with our intuitive convictions about the deontic structure of the moral domain.

To take this idea seriously, however, is to attribute to moral rightness an independent normative significance for deliberation. That we owe it to other people to comply with basic moral standards of conduct is itself something that impinges on deliberation, as a consideration that needs to be taken into account and responded to. Indeed, it impinges on deliberation in the distinctive way I characterized above, as a presumptive requirement (in the generic sense). It is characteristic of the paradigm cases of directed obligation that we register them in practical thought as demands on our agency, specifying things that are “to-be-done” or “not-done” in virtue of the normative nexus that they define. This is, I have contended, the distinctive form of normative significance that directed duties have for the agents who acknowledge them in deliberative thought. The contrast is with other kinds of reason for action, such as considerations of personal enjoyment or convenience, which we register very differently in practical reflection; considerations of this kind speak in favor of our doing certain things, without presuming to exclude from the start the option of acting otherwise.

As Joseph Raz has observed, the difference in these styles of normative relation comes into clear focus in the cases in which they interact. Thus the inconvenience of going out in the rain would ordinarily count in favor of staying home, and it would normally be appropriate to take this consideration into account in just this way in one’s deliberations about what to do. But if I have promised a student I would meet her at the office, the option of remaining at home is off the agenda, and the reasons of convenience that would ordinarily speak in favor of this course of action are rendered inoperative.⁵⁴ Furthermore, the consideration that operates in this way as an exclusionary constraint is the fact that I have undertaken a duty to meet the student in my office, or that I owe it to the student so to act. These are considerations that might be summarized by saying that, in virtue of my promise, meeting the student is now

the right thing to do, and it would be wrong or impermissible to act otherwise.⁵⁵ I conclude that I have to meet the student, because I have in this way obligated myself to do so.

This general approach to morality is at odds with a common conception of moral rightness, as a kind of summary concept that is without independent normative significance.⁵⁶ According to this alternative conception, to judge that something is the right thing to do is to judge, very roughly, that it is what the agent ought to do, all things considered. Likewise, an action will be wrong just in case the agent ought not to do it, taking everything into account. This is the indefinable generic conception of right and wrong that we encountered at the start of section 2.2 above, which marks out actions as “to-be-done” or “not-to-be-done.” Moral rightness and wrongness could then be interpreted, in these terms, as special cases of these generic summary concepts, ones that are in play when there are reasons of a moral nature that speak decisively for or against the agent’s doing something. The result is a kind of buck-passing conception of moral rightness; to say that X is the morally right or wrong action is not to specify a consideration that counts in its own right as a reason for doing X, but to register the existence of other reasons for so acting that are both weighty and moral in character.⁵⁷

Those who accept this kind of account will naturally think that it is fetishistic to be moved to action directly by reflection on the fact that X would be the right thing to do, morally speaking.⁵⁸ The morally conscientious agent is someone who cares about, and responds immediately to, the things that tend to make actions morally right in the buck-passing sense, such as that they would alleviate someone’s need or avoid unnecessary suffering or meet the expectations one has induced another person to form. To care about moral rightness, as something over and above the first-order considerations that make actions right in the first place, is to assign to a consideration the kind of significance for deliberation that it does not intrinsically possess, much as the commodity fetishist invests material objects with the social meanings that are in fact in play only as a result of their being exchanged.

Now one challenge for this summary conception of moral rightness will be to make sense of the deontic structure that seems to be characteristic of the moral realm. If there is something about moral considerations that makes them suited to enter the deliberative field as presumptive constraints on the will, this will not be the fact that the actions they recommend are morally right. That is a consideration that is, by hypothesis, without independent normative significance; as a summary or second-order concept, moral rightness is not

something that we reflect on in deliberating about what to do, but expresses the initial output of deliberation, the conclusion we arrive at by reflecting on the first-order considerations that count as moral reasons in their own right. The burden of accounting for the deontic character of moral considerations will thus fall to the characterization of the reasons that we advert to when we say that an action is morally right.

For all that has been said so far, of course, one of the first-order considerations that contribute to making actions morally right in the summary or generic sense might be the fact that the agent owes it to another party to perform the action, or that the other party has a claim against the agent to such performance. So long as this is just one moral consideration among others, however, we cannot appeal to it to make sense of the idea that moral considerations in general have an inherently peremptory or demanding quality. That is, there will be no unified account of the interpersonal moral domain that itself explains why the considerations that make actions right are in the way of obligations. On the relational interpretation, by contrast, the directed character that we naturally register in moral reflection as a presumptive requirement just is the feature that we are ascribing to actions when we correctly identify them as morally right. It is thus well suited to explain this important dimension of the normativity of the moral domain.

Crucial to this account of moral obligation is the sense of oneself in deliberation as being bound to another person in a nexus of interlocking claims and directed duties. Moral reasons, on the relational conception of them, are part of a normative structure that implicates another party as well as oneself, and this is connected to the characteristic way they impinge on deliberative thought, as presumptive constraints on action and choice. As I put the point above, they belong to a normative nexus that is held in common by two different parties, and one therefore lacks the discretion characteristic of some other normative reasons to discount them unilaterally (for instance, by satisficing rather than choosing the best option, or by acting on a sufficient reason when there is more reason, on balance, to do something else).

To deliberate in these relational terms is not, I think, to fetishize a consideration that is without independent normative significance. It is an important aspect of our social relations that we are connected to other individuals through bonds of directed duty, such as those created by promises and other transactions between individuals. The relational account assumes that there are things that we in this way owe to other individuals quite generally, independently of any transactions we might have entered into with them, and just

in virtue of the fact that they and we inhabit a world together. If this is right, then these general relational duties should likewise have independent significance for practical thought. True, the general relational framework of directed duties and claims is not always salient or explicit in the reflections of the agents who comply with its terms. Agents will in practice often be focused more specifically on the concrete consequences of their actions for the individuals who stand to be affected by them—for instance, on the pain or inconvenience or disappointment that might be caused by their going ahead with a course of action on which they are launched. According to the relational approach, however, thoughts of this kind will implicitly reflect, in morally conscientious agents, an awareness of their significance for directed duties and claims. One will apprehend the interests of the other individual as considerations that potentially ground claims against one to act in one way or another. The structure of one's deliberations, if it were made explicit, would thus be provided by the relational framework.

A different way to develop the idea that obligations constitutively connect us to other parties would be to interpret them in the voluntarist terms sketched earlier, as the commands of an authoritative lawgiver of some kind. According to this approach, thoughts about rightness do not have the same kind of normative significance for the authorities who impose moral requirements that they have for those on whom the requirements are imposed.⁵⁹ It is the act of demanding compliance by the authority that makes certain actions morally wrong in the first place, in the way that involves an obligation not to perform them. In imposing the requirement, then, the relevant authority cannot be understood to be responding to the independent fact that it would be morally wrong in the relevant sense to violate it. Thoughts of moral right and wrong are in this way transparent to the deliberations of the authority whose volitional acts give rise to obligations.⁶⁰ Once the relevant demands have been made, however, thoughts of moral right and wrong are available to figure in the deliberations of the agents on whom the obligations rest. They will see themselves as subject to a presumptive requirement to comply with moral standards of right conduct, just because and insofar as those standards are laid down by the voluntary acts of a suitable authority.

According to the relational alternative, by contrast, obligations are understood by the agents subject to them to constitute a nexus, not with some putative lawgiver whose commands bring them into existence in the first place, but with the specific individuals to whom they are owed. They link agents with persons who have a claim against them to performance. From the perspective

of claimholders, the obligation is not something that comes into effect through a volitional act on their part. It is already there to be cognized and understood, as a requirement that is directed to them, one that goes together with a claim on their part to compliance. It is thus not transparent to the reflections of the claimholders, but instead represents for them, as for the agents, a relational consideration with independent normative significance. Just as promisors understand that they owe it to promisees to perform, so too do promisees understand themselves to have claims against promisors that they so act.

I discuss in the next chapter the distinctive form of normative significance that directed moral obligations have for the individuals to whom they are owed. Before taking up this issue, however, it is important to address a concern that probably already occurred to readers several pages ago. In developing my argument about the bearing of directed duties on deliberation, I have taken for granted our familiarity with this basic normative notion, and focused attention on features of it that suit it to figure in reflection as a basis of practical requirements. We understand ourselves to generate claims in other people when we make promises to them, and the corresponding directed duties make sense to us as considerations that enter deliberation in the guise of presumptive constraints.

But the basic schema of a directed obligation is one that has application in many different domains. There are duties and claims of private law, for instance, which may be distinguished from each other according to the different conventional systems that define them.⁶¹ But there are relational obligations of other kinds as well, including (among others) directed duties and claims defined by the rules of various games. In the sport known in most of the world as football, for instance, there are penalties that attach to various things players might do on the field, and some of these sanctioned actions constitute fouls rather than mere infractions. Many (though not all) fouls, furthermore, have a transparently relational character, insofar as they are recognized to be offenses against an opposing player, generally involving prohibited forms of interference with the player's conduct on the field (such as tackling the player too aggressively, or otherwise impeding the player's free movement to the ball). Within the context of the game, it would be natural for players to take the rules defining such fouls as a sort of relational nexus. Conscientious players will treat these rules as presumptive constraints on their behavior, and actions that flout the rules will be understood to violate specific claims that the players hold against each other in the course of their athletic contest. Outside the context of the game, however, it is not clear that its relational

requirements have much independent significance at all. This raises a natural question about the relational obligations of interpersonal morality: even if we grant that there are such requirements, how can we be confident that they represent important normative considerations, ones that conscientious agents rightly treat as presumptive constraints on their behavior? Perhaps they are no more significant to us than the rules of some random game or social practice that we do not happen to take much interest in.

There are two points to emphasize in response to this natural suggestion. The first has to do with the situations that trigger directed obligations of different kinds. I suggested that in the sports case, there are relational requirements that are partly constitutive of certain games, insofar as they are defined by the rules that make the game what it is in the first place. We might say that the rules of football (for instance) specify what the players owe to each other, just insofar as they are both engaged in a sporting contest of the relevant kind. If this much is correct, however, then it seems that it is one's role as a football player that determines whether it is correct to treat the corresponding relational requirements as presumptive constraints on one's activities. The rules do not even purport to specify what one owes to other people outside the context of the game, and so the question of whether they are important or not, independently of that context, hardly arises. But once one is out on the pitch, playing football against the members of another team, the rules assign claims that the opposing players have against one, and it would reflect a defect in one's reasoning, as a player of the game, not to treat them as constraints on one's own conduct on the field.

According to the relational approach to morality, there are claims that individuals have against us, and corresponding duties that we owe to them, just insofar as they and we are both persons, equally real. Not everyone will accept this contention, of course, and I will address various sources of skepticism about it in the chapters to follow (especially chapter 4). But let us posit, for the time being, that the relational approach is correct in its assignment of claims and directed duties to us, just in our capacity as persons. It seems to follow that we deliberate correctly, insofar as we are persons, only if we treat these person-based duties and claims as presumptive requirements, constraining our decision-making about what to do. Seen in this light, the difference between the requirements of football and moral requirements lies, in part, in the degree to which the situations with which the two systems of relational norms are bound up are escapable. It is always possible for players of football to opt out of the game, and when they do that, the constitutive rules of the

game will no longer be binding on them, as practical requirements. But it is not in the same way possible for us to opt out of the situation in which we are persons whose actions potentially affect other individuals of the same kind. This situation is one that we necessarily inhabit, and so the relational obligations determined by it will figure in the deliberations of all of us as presumptive constraints, if we are well informed and deliberating correctly.⁶²

A second point to emphasize concerns the positive values with which roles of various kinds may be seen to be associated. In the football case, our sense of the importance of the relational requirements constitutive of the game is connected, in part, to our appreciation of the goods that they enable. We admire football players who are able to accomplish impressive feats on the pitch in compliance with the rules of the game they are playing, without fouling their opponents in the process. Those rules impose a discipline on the activities of the players that is closely connected to whatever might be valuable in their role-based performances. This shows itself in the fact that our sense of satisfaction with a player's accomplishments would be diminished by the recognition that those accomplishments came at the expense of an overlooked violation of the rules. (Consider Luis Suárez's uncalled biting foul against Giorgio Chiellini in a decisive group match during the 2014 World Cup tournament in Brazil.)

A different example with the same structure is friendship. It seems to me plausible to suppose that there are things that we owe to other people insofar as we are their friends.⁶³ People who stand in this relationship, and who appreciate its normative significance, will therefore take the claims that their friends have against them as presumptive constraints on their deliberations. But the importance of these requirements is also connected to the values that compliance with them helps to make possible. Friendship is one of the profoundest goods that are achievable in a human life, and our sense of the importance of the demands of friendship is bound up with our acknowledgment of the values that are realized by honoring the claims that our friends have against us. Just as football players enjoy the distinctive values of the game only when they submit to its constitutive rules, so too are the values of friendship available only to those who recognize and honor the requirements that they owe to their friends.

In the moral case, I have not yet said anything to support the idea that there is a characteristic positive value that is enabled by compliance with the relational obligations that we owe to others, just insofar as they are persons. But the idea strikes me as deeply plausible. There is something valuable

about conducting one's life on terms that honor the claims that others have against us. We can take satisfaction in the knowledge that we could look those who might be affected by our actions in the eye, and give an account of ourselves, showing that we lived up to the requirements that we owe to them in particular. We thus *recognize* them as sources of claims against us, and *acknowledge* the significance of those claims in our decisions about the conduct of our own affairs. Furthermore, our sense of the importance of such relational moral requirements is connected to our appreciation of this positive value, which we might refer to as interpersonal recognition.⁶⁴ This is a theme to which I shall return in chapters 3 and 4 below.

In the meantime, some comments about the possibility of conflicts of obligation will bring my discussion in this chapter to a conclusion. I noted above that practical requirements represent presumptive rather than unconditional constraints on the will, in recognition of the fact that circumstances can change in ways that bear normative significance. Emergencies might arise that undermine obligations that would otherwise have obtained, for instance, making it the case that we no longer owe it to someone else to do the very thing that we originally took as a constraint on our behavior. In that scenario, the reasoning internal to morality that supports the original obligation would recognize exceptions. Thus, the interests of promisees that ground entitlements to promissory fidelity are not reasonably understood to generate absolute claims to performance on the promisor's part, come what may. Or it might happen that an obligation remains in force, but comes into conflict with a different obligation whose force is also retained under the circumstances, resulting in a practical dilemma.

According to the relational interpretation, it is possible, for all that has been said so far, that a dilemma of this kind might arise within morality. Moral obligations are directed duties that correspond to claims on the part of individuals against the agent to performance. A conflict within the domain of interpersonal morality in this sense would emerge if there were a situation in which there is no way for an agent to live up to what is owed to one individual without violating the claims that some second individual has against that agent.⁶⁵ But this is at least a coherent possibility. Insofar as the duties in question can be traced to the claims of different individuals, it makes sense that what one needs to do in order to live up to one such claim might preclude the action that is required to honor the other claim. The two obligations would be owed to different persons, each of whom might be entitled to complain if one fails

to fulfill one's obligation to them, even if in doing this one is living up to one's obligation to someone else.

But what about apparent conflicts between morality and other forms of obligation? As I noted above, there are relational requirements of different kinds, associated with discrete normative domains, such as those defined by various private law systems, as well as by conventions, practices, and games of various sorts. I think it is plausible to assume that the relations between many of these systems are ones of hierarchical subordination rather than potential direct conflict.⁶⁶ The intuitive thought here is that directed obligations often function as presumptive constraints in a domain-specific way, and that their force as practical requirements does not extend outside the relevant domain. This prevents them from conflicting directly with the more fundamental requirements that structure and have priority over them. In the game case, for instance, the relational duties of football properly structure one's activities on the pitch, in the situation in which one functions as a player of the game. But this role is normally understood to be subordinate to one's standing as a person who interacts with other individuals, equally real. In a situation in which the relational duties of football appear to conflict with those of morality, it is therefore natural to conclude that the former are no longer fully in force. Thus, if you and I are on opposing teams, and the only way you can prevent me from being hurt by the bomb that has been tossed onto the field is to tackle me from behind without the ball, this is clearly what you owe it to me to do. The fact that you would be fouling me is no longer a pertinent consideration in this context.

A different case that is interesting to consider concerns political obligations. The idea that there are obligations of this kind, such as an obligation to obey the laws of the political communities of which we are members, is not uncontroversial. But to the extent the idea is plausible, the broadly relational model might provide a promising way to understand these obligations.⁶⁷ Political associations involve distinctive forms of interpersonal relationship that can be non-instrumentally valuable, and under these conditions we might well owe it to the other individuals to whom we are so related to comply with the legislative commitments that we have undertaken together. But these obligations, to the extent they obtain at all, are naturally understood to be subordinate to the more basic requirements of morality. We might put this by saying that our membership in a given political community is valuable only insofar as the terms of association for the community meet some threshold of basic moral

acceptability. The result would be that there are independent normative requirements associated with political membership, but not ones that are liable to conflict fundamentally with the requirements of interpersonal morality.

But this does not mean that conflicts between moral and other obligations are impossible. Some directed duties, for instance, are associated with relationships that are fundamental to the meaning of our lives, relationships whose value does not derive from the interpersonal moral values of mutual recognition and regard. There are things that I owe to other individuals insofar as I love them, and stand to them in a relationship of friendship or attachment.⁶⁸ Insofar as these directed duties have a source in values that are independent from those at the heart of morality, they might sometimes diverge from moral requirements. Nor does it seem plausible to insist that the requirements of such special relationships always give way to moral requirements when the two of them come into conflict. There might be a certain sensitivity to moral demands that is built into most reasonable conceptions of friendship and love;⁶⁹ in contrast to the case of political membership, however, it is doubtful that relationships of love and attachment can be valuable only to the extent they are conducted on basically moral terms. So some potential for conflict may go together with our involvement in these relationships.

And there are other possibilities for conflict, as well. I noted in section 2.1 above that there are different models available for understanding the phenomenon of a practical requirement, beyond the relational model that is implicit in the conception of morality as a set of directed obligations. These models, despite their limitations, have at least some application within the terms that they lay out. Thus, a consideration that is anchored firmly in core features of our identity, such that our identity will be threatened or undermined if the consideration is not acted on, is one that it would make sense for us to treat as a presumptive constraint on our planning agency. If my standing as an artist would be betrayed by my failure to take advantage of an important opportunity, this might be something that does not merely count in favor of my so acting, but requires or demands it.

Practical requirements of these various kinds would have sources that are distinct from that of moral obligations, which according to the relational account define what we owe to others just insofar as they are persons who are equally real. But insofar as these independent requirements are also fundamental, the possibility cannot be ruled out that they might sometimes come into intractable conflict with our moral duties. Bernard Williams's probing critical reflections about impersonal morality in its various forms often center

around conflicts of this general kind: between, for instance, an imagined Gauguin's artistic commitments and the moral requirements that are owed to his family, or between our moral responsibilities to strangers and what we owe to the people to whom we are connected by ties of love and affection.⁷⁰

Conflicts of both of these two general kinds—those within morality, involving moral duties owed to different individuals; and those between morality and other sources of practical requirement—raise interesting issues, which are beyond the scope of the present discussion. For now, I would merely observe that nothing in the argument of this chapter precludes the possibility that there are conflicting obligations of either type. On the contrary, the framework I have offered for understanding moral obligation seems to bring with it opportunities for conflicts of both kinds potentially to emerge in the course of living a human life. Insofar as moral obligations represent duties that are owed to individuals, there might be a conflict within morality between duties that are owed to two different parties. And insofar as there is a plurality of models for making sense of something as a practical requirement, the obligations of morality might potentially come into conflict with requirements that have a distinct source, for instance in features of our practical identity or in the demands of the special relationships we stand in to those to whom we are attached.

It is sometimes assumed that a vindication of morality must show that it is not merely a source of obligations, but that its obligations are overriding or supreme, such as to trump any of the normative considerations that might potentially come into conflict with them. It might be nice if there were a philosophical demonstration that this is the case, but it seems to me too much to expect from a theory of moral obligation that it should necessarily secure the supremacy of moral obligations in this ambitious sense. Instead, I have concentrated on a preliminary but to my mind more important task, of explaining why moral considerations should have the distinctive normative force of obligations. As we have seen, considerations of this kind bind the will in the way of practical constraints wherever they remain in force, even if they might sometimes come into conflict with competing practical requirements.⁷¹ The relational interpretation promises to make sense of this aspect of the morality of right and wrong, and this is already an important consideration in its favor.