The Moral Nexus

R. Jay Wallace

Published by Princeton University Press

⇒ For additional information about this book

3

Morality as a Social Phenomenon

IN THE PRECEDING CHAPTER I looked at the normative significance of moral rightness for the agent. I argued, specifically, that moral considerations present themselves as practical requirements of a special kind, and that the relational approach to morality is well positioned to illuminate this dimension of it.

In this chapter, I shift my focus from the agent to those potentially affected by what the agent does. A leading idea here will be that interpersonal morality apparently has normative significance not only for the agent, but for other parties as well, and that it is an important but neglected task for moral theory to make sense of this aspect of it. Moral standards of right and wrong purport to define constraints on agency; but they also purport to provide a basis for interpersonal accountability relations between individuals, articulating what we can expect of each other as each of us pursues our private ends. Disregard of such interpersonal expectations by an agent thus has normative implications for other parties, giving them reason to adjust their attitudes and behavior in response, in the characteristic register of blame.

My main aims in this chapter are as follows. First, I hope to work out the basic idea sketched hastily above, that requirements of moral right and wrong have direct normative significance for other parties besides the agent subject to them. A nondebunking account of them should establish that they are suited by their nature to structure relations of accountability, and that disregard of them by an agent gives others reasons to react in distinctive ways. Second, I shall argue that it is a signal advantage of the relational account of rightness that it provides a satisfying response to this desideratum. According to the relational account, moral rightness is to be understood in terms of directed obligations that are connected to the claims of other parties. But people

who have claims against the agent to compliance with obligations also have reasons to hold the agent accountable for such compliance, and to blame the agent when the obligations are flouted. Finally, I shall identify further elements in our practices of interpersonal accountability that make sense only if the moral norms on which they are based represent relational obligations. In these ways, relational moral obligations are uniquely well suited to provide a normative basis for our social practices of accountability.

The argument of the chapter can be understood as having two complementary phases, which move in different directions between modern conceptions of interpersonal morality and our accountability practices. Phase one, which occupies the first two sections of the chapter, begins with the modern concept of the morally right, arguing that it purports to have normative significance for other parties besides the agent, and that the relational interpretation alone is able to make sense of this dimension of it. Especially important here is the idea that a conception of moral rightness should provide a plausible normative basis for accountability relations, including reactive and other forms of blame. In the second phase, which takes up sections 3.3–3.4, I take a closer look at some of our accountability practices, and argue that they have features that make sense only if the moral norms with which they are connected have a relational structure. Central to this phase of the argument will be the specific reactions characteristic of blame, and the related practices of apology, forgiveness, and moral repair. Taken together, the two phases make the case that interpersonal morality is an inherently social phenomenon, in ways that become fully intelligible only on the relational conception of it.

3.1. The Interpersonal Significance of
Moral Right and Wrong

In the preceding chapter, I suggested that morality seems to have normative significance in the first-person perspective of agency. That a course of action would be morally wrong is not merely an interesting theoretical fact about it, but something that makes demands on the rational will. Indeed, I suggested that moral considerations of this kind present themselves to us not simply as reasons for action, but as obligations, and that it is an important task for an account of morality to make sense of its characteristically deontic form of deliberative significance.

But what is the status of this claim about the apparent normative significance of morality for agents? Does it follow a priori that if X is the right thing

for me to do, then X is something that I am obligated to do? Or is the normative standing of moral considerations as obligations a merely contingent fact about them, one that isn't guaranteed, for instance, by the correct application of the concepts involved in moral judgment? The issues here are delicate ones.

On the one hand, it is certainly part of our modern understanding of morality that it constitutes a domain of practical requirements. Thus we typically cite moral considerations in discussion with other people, as factors that are of peremptory normative significance, requiring or ruling out options that are under active consideration. We raise our children to treat moral considerations in this way, for example. Furthermore, many of us structure our (adult) deliberations on the supposition that moral considerations have this kind of importance, taking facts about rightness and moral value to enter into our deliberations in the way of presumptive constraints. These considerations suggest to me that it belongs to the modern concept of morality that conclusions about what it is right or impermissible to do are imbued with the normative significance of obligations. This is what I meant, in section 2.2 above, in saying that deontic structure is one of the features that characterize the modern concept of morality, a feature that specific conceptions of moral rightness are answerable to. On the other hand, there has to be room for the skeptical position that coherently questions whether people are really obligated to comply with the standards that define what is morally right and wrong. The skeptical position might be mistaken as a matter of fact, but it isn't merely confused; one can grant that it would be wrong to do X, and yet without contradiction deny that this is a reason against acting in that way, never mind a practical requirement.

The best way to do justice to these twin pressures, it seems to me, is to take an element of revisionism to be endemic to the skeptical position. Skeptics, insofar as they deny that moral rightness represents a source of obligations, are asserting the falsity of one of the platitudes that help to fix the modern concept of the morally right.[1] They are thus denying that there is anything in the world that completely answers to this moral concept. It doesn't follow, however, that the position they are adopting is merely confused or incoherent. In saying that people are not in fact required to comply with the standards of moral rightness, they can be interpreted as suggesting that the properties in the world that most closely approximate to our concept of the morally right are not properties that have this kind of normative significance. Morality cannot, as a result, be everything that it represents itself as being, insofar as one

of the platitudes that help to fix the modern concept turns out to be false as a matter of fact. But this strikes me as a plausible thing to say about the kind of skepticism I have been considering. It is a modestly revisionist position, denying something that strikes us as partly constitutive of the modern concept of the moral in the first place, namely the direct significance of the high-level properties it describes as presumptive normative constraints on agency.[2]

I now want to suggest that it is an additional element in the modern understanding of morality that the obligations of moral right and wrong have normative significance for parties other than the agent immediately subject to them. They provide, in particular, a reasonable basis for responsibility relations between individuals, who hold each other accountable for their compliance or lack of compliance with moral obligations. This is a further feature of the modern concept of morality that a substantive conception of it is answerable to. I referred in chapter 1 to the domain of interpersonal morality, suggesting that it collects a set of requirements that is intermediate between the narrow domain of moral rights and the broader class of reasons and values that provide potential bases for assessing exertions of the rational will. But morality in this intermediate sense is interpersonal not merely in virtue of representing obligations that derive from our relations to other people, but also in virtue of collecting considerations that have normative significance *for* other people. A nonskeptical account of morality must make sense of the interpersonal dimension of it, no less than its deontic significance from the first-person standpoint of deliberation about what to do.

There are two closely linked aspects of responsibility that I would particularly wish to emphasize in this connection. The first is the general stance that we adopt toward people when we are prepared to hold them to account for what they do. P. F. Strawson famously adverted to this stance when he contrasted the attitude of "participation or involvement in a human relationship" with a different, objective attitude that may be adopted toward other people.[3] In viewing people objectively, we step back from the expectations that define our ordinary social relationships, and regard them dispassionately, as individuals to be studied or manipulated or cured or understood. Strawson himself did not believe it was possible to sustain an attitude of objectivity in this sense across the board, toward all of the people we interact with as we go about our lives. But it is available to us in the specialized contexts defined by professional activities and roles, such as those of the therapist or the anthropological researcher, and also as an occasional "refuge from the strains of involvement."[4]

The participant attitude, by contrast, is one that we take up toward people by default, and it involves the attribution to them of responsibility for what they do.

Let us call the stance through which responsibility is in this way ascribed to people that of interpersonal accountability. For present purposes, it is not important to take a stand on Strawson's claim that interpersonal accountability relations in this sense represent our default mode of interacting with people.[5] But the contrast with objectivity brings out that accountability has distinctive features that it is open to us to suspend, for strategic or professional or other purposes, at least for limited periods of time. How might these distinctive features best be understood? A promising framework for thinking about them, I believe, is to understand them in terms of demands or expectations.[6] To stand in relations of interpersonal accountability with others is not to be accepting of or indifferent to their conduct, whatever its shape or character, but to hold them to certain standards for comporting themselves. It is to expect or to demand that they should comply with those standards, in a sense that (as Strawson observed) is reasonably familiar to all of us who have been participants in ordinary relationships with other human beings.

To hold people to demands and expectations, in the manner characteristic of accountability, is a way of addressing those demands and expectations to other parties.[7] This stance is connected in turn to distinctive ways of reacting when the demands and expectations at issue are flouted. This brings me to the second aspect of responsibility that I wish to highlight. Those who stand in relations of interpersonal accountability with other people are typically prone to respond to infractions of the relevant expectations with responses that are in the general key of blame. Blame is a complex phenomenon that will probably always elude definitive philosophical analysis; it is simply too protean and multidimensional.[8] I myself favor an approach to it that is indebted to Strawson's work in emphasizing the reactive sentiments, including preeminently resentment, guilt, and indignation, but that can also encompass other, related forms of angry disapproval.[9] Attitudinal reactions of these kinds are paradigmatic and salient examples of blame, even if we wish to concede that blame can take other forms as well.

Reactive blame, as we might call it, insofar as it is emotionally inflected, reflects the fact that the person subject to it has internalized a concern for the expectations or demands that somebody has flouted. When we blame in this way, we reveal that it is not a matter of indifference to us whether agents comply with those standards, but that it matters to us; reactive anger is a way of

being exercised about someone's behavior and attitudes, and it reflects the fact that we are invested in certain basic standards, as ones that govern the person whom we blame. That said, we should allow that blame might take other, less paradigmatically angry forms as well. These extend from expressions of disapproval to comparatively cool modifications of one's relationship with another agent in response to that agent's violation of demands or expectations, such as withdrawal and avoidance.[10] For purposes of discussion in what follows, however, I shall focus primarily on the paradigm form of reactive blame, while also noting the adjustments that might be required in my position in order to accommodate some of the other forms that blame is able to assume.

We have, then, the general stance of interpersonal accountability, whereby we hold people to demands or expectations, as well as specific reactions in the key of blame, to which interpersonal accountability disposes us when those demands and expectations are flouted. The suggestion I would now like to explore is that moral rightness is normative not only for individual agents (insofar as it represents for them a source of practical requirements), but also for these accountability practices, which involve individuals other than the agent.

My basic thought here is that standards of moral rightness are intrinsically suited to figure in an interpersonal practice of accountability. That is, it should make sense, in virtue of the nature of moral obligations, that people hold each other to them, in the way that is characteristically connected with reactive blame. Such obligations thus provide a reasonable basis for social relations that are structured in terms of the addressing of demands or expectations to another party. Furthermore, and relatedly, it should follow from an account of moral rightness that people who act with disregard for the requirements of the moral right thereby give others at least pro tanto reasons for the reactive attitudes involved in blame.

These suggestions reflect a conception of morality as a phenomenon with an essentially social function. Moral obligations, as we have seen, purport to have direct significance for individual agents, providing them with compelling reasons for action that shape their deliberations about what to do. But they are not merely considerations that have this kind of normative importance in the context of first-personal deliberation. They also have an essential role to play in providing the basis for a shared social life, via their role in a practice of interpersonal accountability. This may be thought of as a further element in the characteristically modern concept of morality that I adverted to at the start of the current section.[11]

A reflection of this aspect of interpersonal morality is the thought that we can understand a community's morality by identifying the things that attract reactive attitudes within the community, eliciting responses of angry disapprobation. *Our morality*, in this familiar sense, just is that set of standards that we address to each other in an interpersonal practice of holding one another to account. This is not to say that we are infallible about what morality requires of us. We might attach opprobrium to the wrong things, blaming people when they engage in eccentric sexual practices with other consenting adults, for instance, and failing to get exercised at all when they turn their backs on the basic human needs of vulnerable members our community. We could express this possibility by saying that our moral standards are mistaken or misguided in cases of this kind, and this tells us something about the modern understanding of morality: namely, that we think of morality as a set of standards that function to regulate our responsibility relations with each other, giving people reasons for reactive and other forms of blame when they are violated.

This is an initial statement of my main idea about the interpersonal dimension of the morally right. But the idea requires further clarification and development. I have suggested that there are two ways in which moral rightness seems normative for practices of accountability. First, standards of moral rightness are ones that are intrinsically suited to function as expectations or demands in an interpersonal practice of accountability, and to be addressed by agents to each other within such a practice. Second, violations of these standards provide other parties with reasons for reactive and other forms of blame. But these two aspects are closely connected, and tracing the links between them will help us to understand better the apparent normative significance of moral obligations for the reactions of others.

I said that moral rightness delivers standards that are suited by their nature to define the expectations to which we hold each other within a practice of accountability. This means that it is at least pro tanto reasonable to address those standards to other parties in the ways characteristic of interpersonal accountability. Whether this stance is reasonable on balance arguably depends on ancillary factors, independent of the content or nature of moral rightness, that ultimately bear on the assessment of the stance. For instance, skeptics about moral responsibility argue that it is never ultimately justifiable to hold people to account for their actions, because the conditions of responsible moral agency cannot be satisfied in a world of natural causal processes, or perhaps because the notion of moral responsibility is itself incoherent (so that it would lack application in any possible world we might try to imagine). These

arguments raise very large issues that, though interesting and important, are well beyond the scope of my discussion in this book. Abstracting from global concerns of this kind, however, we can say that standards of moral rightness are ones that it is reasonable to hold people accountable for complying with. That is, assuming that it can ever be reasonable in our world to address demands or expectations to others in a practice of interpersonal accountability, it is reasonable to grant standards of moral rightness this distinctive role. If it would be morally wrong for people to do something—say, breaking their word for reasons of trivial personal convenience—then it is the sort of thing that it is, to that extent at least, reasonable for other parties to expect them not to do.

But when, exactly, is this condition of pro tanto reasonableness satisfied? A plausible answer, in my view, is that it is satisfied, in the first instance, just in case and just because the flouting of the moral standard would give other parties reasons to adopt reactions in the general key of blame. The stance of interpersonal accountability involves a preparedness to blame people when they fail to live up to the expectations to which we hold them in our interactions with them. So standards of moral rightness are suited by their nature to figure in this stance when they are such that their violation gives people reasons for reactive and other forms of blame. It is their potential provision of reasons for blame of this kind that is the primary form of normative significance that moral obligations have for the attitudes and actions of other parties. The view I am developing, then, is that it is part of the modern concept of morality that standards of moral rightness have this kind of significance for blame, and that it is to that extent reasonable to hold others accountable for complying with them.

More precisely, what provide others with reasons to blame an agent are not facts about the violation of moral standards by the agent, but attitudes of disregard on the part of the agent toward the standards. It is possible to act impermissibly, in a suitably objective sense, even if one is reasonably conscientious about one's moral obligations.[12] Having promised to accompany you to the gym to try out one of their Pilates classes, something you are reluctant to do on your own, I will count as acting wrongly if I fail to show up at the appointed time and place. This will be true even if I made scrupulous efforts to see to it that I kept the promise as I planned my activities in the interim. Perhaps I entered the appointment carefully into the Google calendar that I use to keep track of my life, but the entry ends up getting shifted to a different day on account of a malicious cyberattack on the Google servers. Or perhaps, though the calendar does not malfunction, the subway train I board on my

way to the promised appointment does, trapping me short of my destination with no hope of timely egress.

Moral standards define obligations or practical requirements, and as we saw in the previous chapter, agents respond correctly to them when they form intentions to comply with them, treating them as presumptive constraints on their ongoing activities. From the perspective of other parties, then, what it is reasonable to expect or demand of people is only that they will make a conscientious effort to live up to the requirements of the moral right, acknowledging and striving to realize their deliberative role as presumptive constraints. We cannot reasonably expect perfect compliance with the letter of the moral law, insofar as innocent epistemic and other limitations can sometimes lead agents astray even when their attitudinal responses to their moral obligations are beyond reproach. Reasons for blame are provided, accordingly, not by actions of agents that are objectively wrong or impermissible,[13] but by actions that reflect a disregard for the standards of the moral right. It is the flouting of moral obligations—construed broadly, to include the negligent failure to comply with them—that gives others reasons for the reactions characteristic of blame, not their mere violation by the agent. Conversely, if agents make a conscientious effort to comply with moral obligations, this ensures a degree of normative protection from the responsibility-attributing reactions of other parties. Other people might still get angry with agents who act in this way, but such reactions would not be warranted, precisely insofar as the actions to which they respond do not reflect attitudes of disregard for moral standards.[14]

On the general approach to blame that I find most promising, the reasons primarily involved here are reasons for attitudes. The paradigmatic forms of blame are reactions in the general range of angry disapprobation, including most saliently the reactive emotions of resentment, indignation, and guilt. So the reasons that are provided by the disregard of moral requirements are reasons for reactive attitudes of these kinds.[15] But reasons for action may also be in play as well. For one thing, as I noted earlier in this section, there are nonemotional reactions that may legitimately be regarded as forms of blame in at least an extended sense, including the expression of moral disapproval, as well as such actions as avoidance and withdrawal from trust-based interactions with the wrongdoer. For another thing, there are secondary reasons for action associated with the general stance that I have called interpersonal accountability. This stance is, in the first instance, a complex of attitudes through which we can be said to hold other people to expectations, including dispositions to reactive and other forms of blame upon the flouting of those expectations,

and acknowledgment of the legitimate reasons for those attitudes that such deliberate immorality provides. But someone who has these attitudes can reflect on them, and take further measures in regard to them as a result of such reflection.

Our susceptibility to reactive emotions is a refinement of basic psychological mechanisms that function socially to encourage compliance with social norms. We are, as Rousseau was acutely aware, deeply social creatures, and it matters to us profoundly how we are thought of by our fellows, in particular whether people hold attitudes of angry disapprobation toward us on account of what we do. We care about such attitudes not merely because it is disagreeable to experience them, but because we do not wish to inhabit a social world in which such attitudes are harbored toward us, regardless of whether they are expressed to or experienced by us. But our reflective awareness of our participation in this economy of esteem and disesteem constitutes a point of entry for larger practical questions about how we wish to comport ourselves in relation to it, going forward.

If we are moved by the global skepticism about moral responsibility mentioned earlier in this section, for instance, then we might conclude that it is not reasonable on balance to continue to subject individuals to the psychological sanctions that have historically helped to ensure compliance with basic social norms of conduct. (We might be led to this conclusion by the thought that it is unfair to subject people to such sanctions under conditions of general determinism.[16]) If, on the other hand, we are not moved by such skeptical worries, then we may be happy to accede to our participation in the economy of esteem and disesteem through which social norms are traditionally enforced. This would be to acknowledge and act on reasons for affirming the legitimacy of our own tendency to hold others morally to account, and for cultivating a similar disposition in our children, as a way of bringing them to appreciate the importance of social norms and to care about whether they are adhered to.[17]

To sum up, moral rightness appears to be a source of accountability-related reasons of various kinds, including reasons to expect people to comply with the demands that define what it is right and wrong to do; reasons to respond to the disregard of such expectations with reactive and other forms of blame; and reasons to participate knowingly in the economy of social esteem and disesteem that helps to enforce compliance with requirements of the moral right. As I noted, it seems to be part of our modern understanding of interpersonal morality that it has these kinds of normative significance for our

accountability practices, just as it provides individuals with obligations that regulate their deliberations in the way of presumptive constraints on agency.

Note, however, that this conceptual point about interpersonal morality should not be taken to entail that moral standards are in fact reasons for such accountability practices. There is room for a focused skepticism about this aspect of morality that is analogous to the skepticism about moral obligations discussed earlier in this section. This view would hold that the willful or negligent violation of moral standards is not after all something that we have even pro tanto reason to respond to with reactive or other forms of blame. That is, even if we abstract from more global concerns about the concept of responsible agency or about the prospects for its realization in the natural world, we might conclude that attitudes of disregard for moral standards do not in fact provide others with good reasons for angry disapprobation. If I am right, then this debunking position, just like the corresponding skepticism about moral obligations discussed above, would involve some element of revisionism about morality. It would maintain, for instance, that our modern concept of moral rightness or permissibility is not fully realized in the world as we find it, precisely insofar as the standards that determine rightness and permissibility do not appropriately regulate our accountability relations with each other. Our feeling that this outcome would involve a degree of revisionism, however, is a reflection of the basic idea that a normative connection to responsibility relations is built into the modern idea of moral rightness and permissibility.

3.2. Individualistic and Relational Conceptions of the Moral Right

Some moral theories are individualistic; they conceive moral standards primarily as helping to define an ideal of individual achievement, or as articulating the obligations of agents who are not understood to stand in normative relations with other parties.

Consider, for instance, the perfectionist theory of Philippa Foot, discussed in chapter 2, which conceives moral (or better, ethical) standards as means to the realization of an ideal of natural goodness, and as deriving their significance primarily from this function.[18] Such an approach is individualistic, insofar as its locus is the individual agent's relation to an ideal of personal attainment. Whether a given person acts rightly is entirely a question of whether the action reflects traits that human beings need in order to flourish under conditions that are natural to them; it does not in general depend on the claims or

entitlements of other parties. But many other modern theories seem to be individualistic in the same sense. Utilitarianism, for instance, conceives of moral standards in maximizing terms, telling individuals that it is wrong to act in ways that are suboptimal in their effects on the interests of sentient beings. Utilitarian rightness depends in part on how other individuals react causally to the things that an agent might do, but it does not depend on any normative relations that the agent stands in to those individuals.

Individualistic approaches of this kind appear to leave it mysterious why attitudes of disregard for the standards of the moral right should provide others with a normative basis for relations of interpersonal accountability. For example, utilitarians often hold that lifestyles of bourgeois consumption are morally objectionable in a world in which there are vast numbers of people living in conditions of extreme need. So long as these conditions persist, individual expenditures on consumer goods will be hard to justify in utilitarian terms, insofar as greater utility would be achieved by donating the funds at issue to an organization such as Doctors without Borders or GiveDirectly. But it is another matter entirely whether we should blame individuals when they display attitudes of disregard toward this moral consideration.

On the utilitarian approach, blame should be treated like any other intervention into the causal order, and assessed by reference to its likely effects on the welfare of those affected by it. Thus, it is sometimes argued that it would be wrong to blame affluent individuals when they fail to organize their personal lives according to the principles of utilitarian consumption. Doing so might just discourage and demoralize them, making it clear how onerous morality would be if they actually took it seriously as a basis for ordering their lives, with the effect that those who are blamed would do even less to contribute to improving the conditions of the many millions of people whose lives are characterized by deprivation and disease.[19] For these same reasons, it might conceivably be for the best not to encourage children to internalize the kind of emotional commitment to utilitarian standards that would render them disposed to react to deliberate moral infractions with such sentiments as resentment and indignation.[20] The question of the attitudes of an individual toward moral standards seems to be completely independent from the question of whether those same attitudes provide others with a normative basis for interpersonal accountability relations.[21]

Against this, it might be said that it is open to utilitarians to distinguish more carefully between the attitudinal and the behavioral dimensions of blame. Actions that serve to sanction people for their moral lapses are certainly, on the

utilitarian view, to be assessed with an eye to their consequences for human welfare. But the reactive attitudes that register blame emotionally, such as resentment and indignation, might be taken to function differently, answering to considerations that render them intrinsically warranted or fitting. I do not believe that this move will save the utilitarian, however. As I shall argue in more detail in sections 3.3 and 3.4 below, the considerations that we understand to provide reasons for attitudinal blame are not captured by individualistic conceptions of the right such as utilitarianism. Individuals who fall short of the conduct prescribed by the principle of utility, or who do not strive to act in accordance with that principle, may be going astray relative to normative standards that in some sense apply to them. But it is not at all clear why this should be thought to render warranted focused reactive attitudes such as resentment or indignation on the part of other people, any more than it provides a basis for the behavioral expressions of blame.

Individualist approaches such as utilitarianism thus seem to deny the element in the modern concept of morality that I have been at pains to explicate in this chapter. According to that concept, standards of moral right and wrong are such that the agent's attitudes toward them have direct normative significance for the responsibility reactions of others. Moral standards must be ones that it is reasonable to hold people accountable for complying with, in the ways that characteristically involve a susceptibility to blame. There must be something about them that gives an agent's attitudes toward them a special normative significance for our responsibility reactions, protecting from opprobrium those who conscientiously strive to comply with them, and rendering those who are indifferent to them specially vulnerable to reactive and other forms of blame. Utilitarian and perfectionistic theories seem deficient when viewed in this light. They take moral requirements to derive from an individual's relation to impersonal value or to an ideal of human attainment. But there is nothing in the nature of such requirements that would seem to explain why our attitudes toward them necessarily have normative significance for accountability relations. Why should other people become exercised by the fact that you are indifferent to standards of human perfection in the conduct of your affairs, or that you do not fully live up to the rigorous demands made on you by an ideal of impersonal value? An individualistic theory might turn out to be correct at the end of the day. If the argument of the preceding section is sound, however, it will follow that such theories are revisionist in at least some measure, denying an interpersonal dimension of morality that seems to us to be essential to it.[22]

In response, it might be suggested that there are individualistic theories that do not leave it a contingent matter whether we have reason to respond to deliberate wrongdoing with reactive and other forms of blame. Consider a version of perfectionism that holds that it is a virtue, part of the ideal of a flourishing human life, for agents to internalize the kind of concern for moral standards that makes them prone to blaming reactions when people flout those standards. On a theory of this kind, morality itself would enjoin us to respond with something like blame to instances of deliberate wrongdoing, and it would therefore seem that there is a nonaccidental connection between such behavior and the responsibility reactions.

A theory of this kind would not really do justice, however, to the aspect of morality described in section 3.1. There are two aspects to the problem. First, on the perfectionist account, what makes actions wrong, in general terms, is that performing them would be incompatible with virtue. This is an individualistic approach, as I noted above, which holds that the rightness and wrongness of actions are matters of the agent's relation to an ideal of human perfection. But there is nothing in this general way of thinking about morality that explains why an agent's attitudes toward rightness and wrongness should *themselves* provide a normative basis for the responsibility reactions. It is only when we supplement the general account with a specific, substantive theory of virtue that we introduce a reason for people to adopt blaming responses toward episodes of deliberate wrongdoing.

Second, the reason that is introduced by the substantive theory of virtue seems to be a reason of the wrong kind. I should react to deliberate wrongdoing on the part of others, according to the perfectionist theory, because a failure to do so would instantiate a vice on my part. My reason for blaming, in other words, is provided by considerations having to do with my own relation to an ideal of human attainment.[23] According to the view I have been developing, by contrast, what provides me with a reason for the responsibility reactions is, in the first instance, the attitudes of the person whom I would blame. It is because that person has acted with indifference or contempt toward moral standards that I have reason to hold the person responsible.[24] This is precisely the normative connection between rightness and responsibility that I have been developing, and it points toward an essentially interpersonal dimension of morality. Insofar as the perfectionist theory under consideration remains individualistic, it continues to have a revisionistic character, even if it yields the conclusion that morality gives us a reason of some kind to blame others when they flout moral standards.[25]

Consider next the divine command theory.[26] This approach, like other versions of voluntarism, is not individualistic, insofar as it understands moral rightness by appeal to the normative relations that agents stand in to figures who have authority over them. But its non-individualistic character does not shed light on the aspect of morality that is currently under investigation. According to the divine command theory, there may be things that it would be morally good to do, independently of whether we are commanded by God to do them. But we stand under obligations to act in these ways only in virtue of the fact that God lays down a requirement on us so to act. God is thus the source of the moral obligations that are binding on us, and we owe it to God to act in accordance with those obligations.[27] As we saw in the preceding chapter, an approach along these lines might go some way toward explaining the significance of obligations in the first-personal perspective of deliberation, their standing as genuine requirements that make claims on the agent's will. But it does not offer an illuminating account of the social dimension of moral obligations, where this is taken to involve their suitability to structure an interpersonal practice of mutual accountability.

To stand in relations of accountability to another person, or to address moral demands to the person, is, on the account sketched in section 3.1, to hold the person to the demands in the ways constitutively connected to broadly reactive attitudes. So the question is whether it is reasonable to hold people to demands in this reactive way if those demands derive from the commands of a benevolent and omnipotent deity. I think there is a real question whether it is. The worry, in a nutshell, is that it isn't really anyone's business whether other people live up to the requirements that are owed by them, individually, to God. It is between them and the deity, as we might put it, and hence not something that the rest of us have good reason to get exercised about, in the ways characteristic of interpersonal accountability.

Granted, many of the divine commandments may prescribe behaviors that affect other persons in our local or less local community, insofar as they include demands to act charitably or with justice (for instance). Those who are members of our social world are in this way apt to be affected, in one way or another, by our compliance or lack of compliance with the requirements that God imposes on us. But this is completely incidental to the status of the requirements as genuine obligations. As was noted in the preceding chapter, God could in the very same sense lay down requirements on us that have nothing to do with our relations to other people, such as dietary restrictions on the consumption of flora and fauna of various kinds. The account of moral

rightness on offer thus leaves it mysterious why things that are morally required are for that very reason also things that it is reasonable to address to the agents whom the requirements regulate, in the way characteristic of a practice of mutual accountability.

As with the individualist theories considered earlier, it would be possible to make substantive additions to the divine command theory that could help shed light on why it is reasonable for people to hold one another to the moral requirements commanded by God. For instance, we might, once again, appeal to a specific theory of virtue, one that has it that virtuous agents adopt the stance of holding each other to moral requirements, as a basis for relations of accountability. This might be a good disposition for individuals to cultivate and acquire, for instance, insofar as its presence in a community helps the members of the community to live up to the demands that God has addressed to them individually. Or perhaps it is a further divine commandment that each of us should cultivate a disposition to reactive moral address with regard to the other requirements that God has imposed on us; we would then owe it to God, so to speak, to treat those requirements as a basis for a practice of mutual accountability.

But these would be ancillary stipulations, going beyond the account of moral rightness itself, which identify reasons for holding people to moral obligations that are potentially of the wrong kind. The stipulations are no more successful within this framework than in the case of the individualistic theories canvased above. The core account of moral obligations latent in this version of voluntarism traces them to the relation that individuals stand in to the deity, insofar as they are addressees of divine commandments. Nothing in this core account, however, explains why moral obligations, so understood, should also function as the basis of a practice of essentially interpersonal accountability.

Let us turn, finally, to a different non-individualist approach, namely the relational account that is my main topic in this book. According to the relational approach, moral rightness essentially defines a set of requirements that are owed to other individuals in particular, where those individuals have claims against the agent to compliance with the requirements that are directed to them. I now wish to argue that an account of this kind is peculiarly well suited to make sense of the social aspect of moral requirements, their status as considerations that appropriately structure relations of mutual accountability.

Thus, suppose that you owe it to me morally to keep the promise you made to me, or to refrain from bodily trespass as you encounter me in the street. According to the relational account, these requirements essentially concern

our relations to each other, and they are connected constitutively to claims that I have against you to compliance with them. Insofar as I have a claim against you to compliance, however, it would seem perfectly reasonable for me to hold you to the requirements in the way that we have seen to be characteristic of interpersonal accountability. Indeed, moral requirements of this kind seem tailor-made, in virtue of their relational structure, to be requirements that function as a basis for a practice of accountability to other parties.

Note, for one thing, that it is a noncontingent feature of relational requirements of this kind that they have significance for another party. They are owed to other individuals, to whom the agent is, as Gary Watson has put it, "beholden."[28] But if a requirement that you are under is in this way directed to me, then it is necessarily not a matter of indifference to me whether you make a conscientious effort to comply with it. Rather it is my business, in virtue of the relational structure that makes the requirement obligatory in the first place. As a claimholder, it is thus fully reasonable that I should take an interest in whether the requirements that are owed to me are lived up to and taken seriously by those who stand under them. And holding someone to the requirement, in the spirit of interpersonal accountability, is the natural way of doing this.

Consider, next, what happens if a directed moral requirement is flouted. The agent who failed to take it seriously will have shown a certain disregard for the values at the heart of morality. But moral obligations, on the relational account of them, are essentially connected to individual claims; so disregard for such requirements is eo ipso disregard for the person to whom the requirements are owed. There are, more specifically, two distinct forms that such disregard for persons could take. First, agents who flout a directed requirement could acknowledge the persons to whom it is owed as the bearers of claims, but fail to take the claims seriously in their deliberations about what to do. Second, the immoral agents could treat the persons to whom the requirement is owed as if they didn't hold claims to performance in the first place. In the first case, the agents are disregarding the other parties by neglecting to honor the claims the parties hold against them; in the second case, disregard of the other parties takes the more basic form of failing to take seriously their standing as bearers of claims.

But showing disregard for other parties, in either of these ways, is precisely the kind of stance that would seem to provide a reason for the characteristic reactions of blame. In failing to take the moral requirement seriously, you have

not merely fallen short of some personal code of conduct that pertains to you alone (or to your relation to the divinity). You have displayed indifference to or contempt for another person, as someone whose interests undergird claims the latter holds against you, and this is the primary normative basis of the reactive attitudes. Thus resentment, which we may take to be a paradigm of reactive blame, is constitutively a response to the fact not just that something bad has happened in a person's immediate environs, but that—to recall an expression introduced in chapter 1—the subject of the attitude has suffered a specifically moral injury through the actions of another party. It in this way presupposes a relational understanding of the requirements of the moral right that give it content.

The point here is a deceptively simple one, but it is of great significance for understanding the connection between moral rightness and interpersonal accountability. It is something of a commonplace in the literature on responsibility to characterize blame, at least in its paradigmatic forms, as a response not merely to actions that are morally wrong, but to actions that wrong another party. Thus, Miranda Fricker writes of what she calls "communicative blame" that it is "a cognitively loaded moral emotion," whose content involves "a perception of a wrong one suffers at the hands of another."[29] Similarly, Gary Watson writes of moral norms that "those who violate these norms wrong others, which explains the distinctive significance for interpersonal relations these standards possess."[30] Another example is provided by Pamela Hieronymi, who observes that "your resentment of my impatience marks the fact that you have been wronged by someone, the quality of whose will matters."[31]

These remarks are offered by the quoted philosophers without much explicit defense; they are apparently regarded as sufficiently plausible that it can be taken for granted that the reader will assent to them as soon as they are formulated. But it is worth dwelling at least briefly on these ways of speaking, for they point toward two connected truths about responsibility and blame that are of profound importance for the argument of this chapter. First, the characteristic occasion for blame is the fact that one party has acted in a way that wrongs another; this is the circumstance that we intuitively understand to provide people with reasons for reactive and other forms of blame. Second, those reactions themselves incorporate an appreciation of the fact that their target has done something to wrong another party. Reactive blame is cognitively structured, and its content is naturally cashed out in terms of essentially relational norms; as Fricker puts it, resentment incorporates "a perception of

a wrong one suffers at the hands of another." It is not merely the fact that agents have acted wrongly, but that they have wronged somebody, that is both the reason for blame and its cognitive focus.[32]

Note, however, that both of these points presuppose something that was emphasized in section 1.1 above. This is that the relational conception of a wrong or a moral injury is an attitude-dependent one. The kind of wrong or injury that constitutes an occasion for blame, and that figures in the content of accountability-ascribing attitudes, is not an objective violation of a moral obligation, but involves an attitude of disregard for the obligation on the agent's part. To return to the example from section 3.1 above, I do not wrong you if I fail to show up to the Pilates class on account of a malicious attack on the server that maintains my calendar or a breakdown on the train I am taking to get to the class. In cases such as these, I might be said to have violated the duty I owed to you, but I didn't flout the duty, and it is only in cases that involve something like flouting or neglecting a duty that I can be said to have wronged you, or caused you a moral injury. But in fact, this is precisely the kind of connection between obligation and reasons for blame that we were trying to make sense of. We wished to understand how an attitude of disregard for moral requirements could itself constitute a pro tanto reason for reactive and other forms of blame. The answer is that it does this if the requirements themselves are connected to claims held against the agent by another party, as they essentially are on the relational account.

A similar line of thought is implicit, I think, in Strawson's highly influential account of our interpersonal accountability practices. As noted in the preceding section, Strawson favors an approach to these practices that emphasizes the reactive attitudes, a paradigm of which is resentment. He notes, further, that the occasion for resentment is an "offense" or "injury" to one party by another,[33] and that offenses or injuries in the relevant sense are in turn connected to the "quality of the will" of the offending party. Thus the reactive attitudes "are essentially natural human reactions to the good will or ill will or indifference of others towards us, as displayed in *their* attitudes and actions."[34] Finally, Strawson emphasizes the connection between these reactions to ill will and the notion of a demand, observing that such sentiments as resentment reflect demands or expectations to which we hold people in our ongoing interactions with them.[35] Let us call these three elements in his position injury, quality of will, and demand.

Strawson does not himself advocate explicitly for a relational interpretation of the demands or expectations characteristic of the participant standpoint.

But it seems to me that a relational account of this kind represents the natural tendency of his position; it is the interpretation that makes best sense of the connections he posits between the elements of injury, quality of will, and demand. Thus, Strawson assumes that the demands to which we hold each other in practice are such that disregard of them will itself involve an attitude of disregard for or indifference toward another person, constituting thereby an offense or injury against that person. As we have seen, however, the relational interpretation of demands is best equipped to make sense of these connections. On that interpretation, moral demands are constitutively connected to claims that other individuals have against the agent to compliance with them. So disregard for relational demands will essentially involve disregard for those individuals (in the form, specifically, either of disregard for their claims, or of disregard for their standing as bearers of claims). This in turn is the sort of "quality of will" that we intuitively understand to constitute a kind of offense or moral injury against claimholders, providing them with a normative basis for resentment.[36]

The ideas that I have been developing in this section, concerning the normative significance for other parties of the violation of relational moral requirements, are connected to the positive value that is enabled through compliance with such requirements. In section 2.4 above, I referred in this connection to the value of interpersonal recognition, which is realized when we conduct ourselves on terms that fully acknowledge and do justice to the claims that other parties hold against us, simply as persons. When we satisfy this condition, we will be in a position to give an account of ourselves to each of the individuals who may have been affected by our conduct, one that shows us to have honored the obligations that are owed to them in particular. Interpersonal recognition, construed in this way, is the converse of the situation we find ourselves in when we flout or neglect the requirements of relational morality. Disregarding such requirements, I have suggested, amounts to disregarding the claims that someone else holds against us, which in turn is a form of disregard for them as a person. By the same token, however, conscientious regard for the directed obligations that are owed to others is a way of acknowledging their claims against one, and so displaying recognition for their standing as morally significant persons.

I have argued, further, that disregarding the claims of others renders one vulnerable to the negative reactions characteristic of blame, providing others with a normative basis for these reactions. If this is right, however, then instantiating the positive value of interpersonal recognition in one's relations

to others should ensure a measure of protection from these unwelcome reactions. Of course, other parties might subject us to opprobrium and blame even when we have not done anything that would genuinely warrant these responses. People are sometimes irrational, petty, and vindictive, and nothing we might do can guarantee that we will never be the target of negative attitudes that are untethered from any normative basis in the facts about our relations to them. But there is a melancholic satisfaction to be taken, even here, in the knowledge that these reactions are not in fact fitting responses to our behavior.

Interpersonal recognition is achieved when we act in a way that deprives others of a warrant to resent our treatment of them. In acknowledging the significance of their claims against us, as constraints on our decision-making, we also acknowledge their moral standing as individuals whose interests matter equally. Doing this does not ensure that we will live in harmony with the actual individuals who share our social world, but it represents a valuable contribution to the normative basis of such a community.[37]

3.3. The Relational Structure of Interpersonal Accountability

Proponents of utilitarianism (or other consequentialist views) may be tempted to respond to the argument of the preceding section that their theory, too, can make sense of the role of moral requirements as bases for a practice of interpersonal accountability. The crux, as we have seen, is to explain why disregard for moral requirements should itself provide other parties with reasons for reactive and other forms of blame. The utilitarian understands the standards of the moral right to require that we maximize the impartial good, taking equally into account the interests of all those potentially affected by our actions. But then it might seem that to disregard moral requirements, so understood, is to fail to regard the interests of some person or persons as worthy of equal consideration. And this failure of regard, it might be argued, is something that could provide a normative basis for the blaming reactions that are characteristic of interpersonal accountability.

The inference involved in this argument, however, seems questionable. One potential problem is that it appears possible to disregard standards of consequentialist rightness without formally denying the equal standing of those affected by our actions. Climate change skeptics, for instance, might insist that there is no moral requirement for members of affluent societies to reduce the

carbon emissions associated with their activities, not because they deny that peoples' interests ultimately count equally, but because they deny the consensus of scientists about human contributions to whatever changes in global temperature might be occurring. Consequentialist rightness requires individuals to maximize the impartial good, and this involves both a willingness to treat peoples' interests equally, and a willingness to reason soundly about the causal effects of the various actions that it is open to one to perform.

Against this, it might be replied that one does not really count as having disregarded the consequentialist standards if one's failure to meet their terms results from nonculpable forms of mistaken causal reasoning. In the preceding section, I suggested that disregard for relational moral requirements involves a failure to make a conscientious effort to comply with them, or a failure to take them seriously in one's deliberations. But there need be no failure of these kinds, in relation to consequentialist standards, if agents who fall short of them do so through innocent mistakes in reasoning about the effects of the actions that it is open to them to perform.[38]

There is a different and more serious difficulty for the consequentialist account of the social dimension of morality, which was mentioned briefly in section 3.2 above. Consequentialists will ultimately condone actions through which accountability is interpersonally expressed only if those actions contribute to the overall maximization of the good. As we saw, however, there are many cases in which it seems dubious that it would be for the best to subject individuals to blame when they fall short of the standards of the consequentialist right. Consider the many people who continue to spend discretionary resources on indulgences for themselves and their friends and family members, rather than contributing those resources to alleviating the plight of strangers who are much worse off than they are. Such agents might be failing to count the interests of all people equally, in the way consequentialist rightness would require. But it would arguably be suboptimal to blame them under these circumstances, since that would lead them to dismiss moral requirements altogether rather than to do what they can to comply with them.

In response, the consequentialist might maintain that disregard of the relevant standards of right conduct is at least a pro tanto reason for reactive and other forms of blame, even if it might not be warranted on balance to subject agents to those reactions whenever they flout consequentialist requirements. The pro tanto reason for blame, it might be said, just is the failure of equal consideration that is implicit when agents disregard, in the relevant sense, consequentialist standards of rightness. But even this residual connection between

consequentialist rightness and accountability seems to me questionable. As I noted in section 2.2 above, consequentialism offers a particular interpretation of the cosmopolitan ideal latent in modern approaches to morality, according to which we operate as agents in a world that includes others who are equally real, and whose interests matter equally for moral thought. The consequentialist interpretation of rightness tells us to attach equal weight to the interests of each person in maximizing the good, and it is an appealing feature of that interpretation that it gives expression to the more fundamental ideal of the equal standing of persons within moral reflection.

Note, however, that it is no part of the consequentialist conception that individuals have *claims* against us, as agents, that we should in this way attach equal weight to their interests. Consequentialism is an individualistic account of rightness, in the terms of the preceding section, which articulates requirements of right conduct in abstraction from assumptions about the normative relations that the agent stands in to other parties. Within this individualist context, however, it is not at all clear that a failure to attach equal weight to someone's interests is an attitude that provides that person with even a pro tanto reason for reactive and other forms of blame. Such a failure may represent a criticizable defect on the agent's part, relative to an applicable norm of individual attainment, one that involves a flaw in reasoning about the interests of another person. But it is not something that could be said to *wrong* the person whose interests are at issue, or otherwise to constitute something like a moral injury to that person.

Insofar as we intuitively think of a failure to attach equal significance to someone's interests as a basis for blame, I submit that we are taking for granted the relational framework that was shown above to be implicit in our ordinary understanding of the circumstances that warrant blame. That is, not only are we assuming that individuals are required to attach equal importance to the interests of others, by attempting to maximize the impartial good in their actions (something that is controversial in itself, as an interpretation of the more basic ideal of the equal moral standing of persons). We are assuming, more specifically, that they owe it to each other to deliberate about action in this way. This is the framework of normative relations that would make the disregard of consequentialist standards something that provides a pro tanto normative basis for accountability reactions. But this relational framework is foreign to the consequentialist approach to moral rightness.

These remarks reinforce the point made in the preceding section, which is that we commonly (if implicitly) understand the circumstances that occasion

reactive blame in relational terms, as actions that wrong another party. I now want to argue that there are additional features of our practices of interpersonal accountability that are essentially positional, and that their positional aspect equally presupposes a relational account of the requirements that underlie those practices. These features are apology, forgiveness, and repair; let us take them up in turn.

It is part of our common moral understanding that wrongdoing calls for a distinctive kind of response on the part of the agent who flouts moral requirements. Specifically, it calls for acknowledgment of wrongdoing and apology. In saying that these responses are called for, I am saying, in part, that there will be a compelling basis for others to blame the agent at least as long as these responses are not forthcoming. Just as wrongdoing provides a normative basis for reactive and other forms of blame, so too do these reactions give agents a reason to acknowledge their wrongdoing and to apologize for their trespass against moral requirements. Indeed, it is not a stretch to suppose that responses of this kind are morally obligatory on the part of agents who have acted impermissibly. To think this is the case is to suppose that there is a secondary obligation that applies when people have flouted a primary moral obligation, the content of which is that the agents who have acted wrongly should acknowledge their moral failing and apologize for what they have done.[39]

It is a striking fact about our accountability practices, however, that our understanding of these secondary obligations is itself relational. A natural way to express the secondary obligations would be to say that wrongdoers owe an apology to those they have wronged through their actions. That is, there is another party who is uniquely positioned to be the recipient of the agent's acknowledgment of wrongdoing and apology, and the duty to respond in these ways is directed to that party in particular. Thus, suppose that Nora and Sebastian are playing with Legos, and that in a fit of envy and frustration Nora smashes the elaborate castle that Sebastian has been working on for the past half hour. Nora might confess what she has done to the supervising parental authority, and say that she is sorry for her action. But it would be perfectly natural for the parent to respond by telling Nora that she needs to apologize to *Sebastian* for destroying the structure he was working on (and to add, perhaps, that she needs to do this *nicely*). That is, we induct people into our accountability practices by encouraging them to appreciate that these agential responses to wrongdoing are owed to another party in particular. They cannot be discharged by directing an apology to just anyone.

Furthermore, it is equally implicit in those practices that the party to whom acknowledgment and apology are in this way owed is the person who was originally wronged by the impermissible actions that now call for apology. Nora owes Sebastian an apology for her willful action, because it was Sebastian who was wronged by that action in the first place. He had a claim against her not to treat him in that way, which was flouted by what she did. Apology and acknowledgment are thus doubly relational. Not only do we understand the secondary obligations of apology to be directed to another party; our understanding of their direction presupposes that the primary obligations whose violation gives rise to them are themselves directed to the same party. There is a positional aspect to our practices of apology and acknowledgment that makes sense only against the background of a deeply relational interpretation of basic moral requirements.

Consider next forgiveness. This is a stance that is normally adopted after the point in time at which the agent of wrongdoing has apologized and acknowledged moral fault. It is a complicated business, involving on the one hand the tacit recognition that the agent acted in a way that makes reactive blame fitting, but forswearing, on the other hand, the reaction that would in this way be warranted. There are difficult questions raised by forgiveness about how these different aspects of it can be combined, and also about the conditions that paradigmatically occasion it and the purposes that it serves.[40] I would like to set these interesting questions to side for the moment, however, in order to focus on a more elementary point about forgiveness. This is that we commonly understand it to be, like apology, a profoundly positional phenomenon.

Forgiveness may be understood to involve the exercise of a kind of normative authority. In bestowing forgiveness, we step back reflectively from the reactive blame that we understand to be rendered fitting by what an agent has done to us, and disavow that reaction, going forward, as a significant factor in our continuing relations with the agent. To forgive is not necessarily to overcome all resentment in one's feelings about the wrongdoer, but to adopt a distinctive stance toward such attitudes, forswearing them as attitudes that should be accorded importance in one's ongoing relations to the agent of wrongful action. The individual who bestows forgiveness is someone who is authorized to adopt this kind of reflective stance, and who has a measure of discretion about when and in what ways to make use of this authority. Thus, even after the wrongdoer has acknowledged moral fault and issued a remorseful apology, it is not the case that forgiveness can be demanded as a matter of right or entitlement. Apology and acknowledgment of fault do not, after all,

undo the fact of moral injury, which continues to represent a consideration that would render reactive blame fitting or warranted. Those who are authorized to bestow forgiveness thus have some leeway to decide for themselves whether and under what circumstances it will be forthcoming.

Note, however, that the authority that is thus presupposed by forgiveness is the proprietary possession of only some individuals, who stand in a distinctive relation to the wrongdoing that might be forgiven. Thus, it is not the case that just anyone is empowered to forgive a wrongdoer's moral transgressions. Just as apologies are owed specifically to those who have been wronged by what the agent did, so too are those individuals alone authorized to bestow forgiveness on the wrongdoer.[41] Once again, we see that there is a positional aspect to our practices of interpersonal accountability, which presupposes that the norms that structure those practices are relational in their content. Third parties lack the authority to bestow forgiveness for actions that flout moral requirements, precisely insofar as they have not themselves been wronged by those actions. The authority to forgive is a retrospective residue of the claim that was disregarded through the wrongdoer's action. It is only those to whom compliance with moral requirements was originally owed who are empowered to forgive wrongdoers for their transgressions.

Indeed, it seems to me that forgiveness would not really make sense, as a distinctive practice, in the context of an individualistic interpretation of moral obligations. One might attempt to emulate that practice within an individualistic framework by stipulating (say) that those individuals should be viewed as authorized to bestow forgiveness who have been harmed, in a nonmoral sense, through actions that flout nonrelational standards of conduct. The positional element in forgiveness would be operationalized, according to this interpretation, through the causal relation between an action that is individualistically defective and harm caused by that action to another party (or parties). But the resulting account seems implausible.

I have suggested that the power to forgive is a form of normative authority or entitlement, which the individuals who are so empowered have discretion to exercise or not, as they see fit. But it is mysterious why standing in a merely causal relation of harm to an act that is personally deficient should ground a normative status of this kind. The point is not the metaethical one that normative statuses cannot be based or grounded in nonnormative facts, for at the end of the day, even a robustly realist conception of the normative should allow that there are systematic relations between the normative and the nonnormative.[42] The point, rather, is that the individualistic account does not

identify a convincing basis of any kind for the normative status involved in forgiveness. The causal connection between individually deficient conduct and harm to another party isolates a contingent effect of wrongdoing, rather than something that is essential to it. But our conception of people as having authority to forgive is not hostage to whether they happen to have been harmed on balance by the wrongful conduct of another party. Rather, it is of a piece with our sense that they have suffered a distinctively moral injury through such conduct, where this in turn presupposes that they have a claim against the party not to have acted in that way. As I put it above, the authority to forgive is a residue of the claim that was originally flouted through the wrongful action that is a candidate for forgiveness.[43] Severed from the context of this kind of normative relation between the agent of wrongful action and the person whom it wrongs, it seems that the essentially positional notion of the authority to forgive would be out of place.

Consider, finally, the phenomenon of moral repair. It is customary to think that the person who transgresses against moral requirements is under a duty not only to acknowledge wrongdoing and to offer an apology, but also to do what is possible to make amends. Like apology and forgiveness, this element in our accountability practices is deeply positional. The obligation that the wrongdoer is under to make amends is owed to one individual (or set of individuals) in particular; this is, as in the case of apology, the individual who was wronged by the agent's action in the first place. We thus encounter, once again, a structure that is doubly relational, with a directed secondary obligation that is parasitic on our understanding that the action that gave rise to it involved the flouting of a primary obligation that was likewise relational.

Here, more plausibly perhaps than in the cases of apology and forgiveness, we can offer a nonrelational reconstruction of our practice that at least makes partial sense of it. In this spirit, it might be suggested that the duty to make amends can be explained as an equitable way of allocating the costs incurred through an individual act of wrongdoing. Insofar as such acts cause adverse effects for other persons, there is a need to make up for them, so far as it is possible to do so, and the costs of satisfying this need seem most fairly assigned to the individuals whose actions gave rise to it in the first place.

But this quasi-economic way of thinking about repair neglects an important dimension of our accountability practices. Those who flout relational moral requirements show disregard for the individuals to whom they are owed; they thereby impair the relationships they stand in to the persons whom they

wrong, and this is a reasonable basis for the secondary obligation to make amends that is also directed to those individuals. The quasi-economic interpretation of repair, by contrast, does not really make compelling sense of the directed character of the duty that it defines. It may be equitable for the social cost of redress to be assigned to the agents whose voluntary violation of individualistic requirements caused the damage that now requires compensation. But it does not follow from this observation that the duty to provide redress should be owed by the agents to the parties who were harmed. To render intelligible this aspect of the practice of repair, we need to embed it within an essentially relational framework, tracing the original wrongs that give rise to the duty of repair to the disregard of the claims held by other individuals against the agents who acted wrongly. It is because those agents have impaired their relationships to the claimholders, through actions that wrong those persons in particular, that the secondary obligation of repair is owed to the same parties. Once again, the secondary duties can be understood as retrospective residues of the claims that were originally flouted.

This way of speaking is reminiscent of a formulation that Arthur Ripstein has used to characterize the duty of remedy in the private law of torts. That duty, Ripstein says, needs to be understood to reflect the fundamental idea that the right or entitlement flouted through the tortious action "survives its own violation."[44] The duty of redress on the part of the tortfeasor should be conceptualized in relation to this persisting entitlement, as a duty owed to the plaintiff to enforce the entitlement that lives on after the original tortious action. But this characterization of the normative situation, while suggestive, seems to me ultimately somewhat misleading, both in application to the private law of torts and in extended application to the generalized moral case.

On the one hand, there is a perfectly trivial sense in which an entitlement or claim, as a normative notion, survives its own violation. That is, the fact that the claim was not honored does not entail or bring it about that the claim was not after all in place; this shows itself in the fact that we understand the wrongful action, retrospectively, precisely to be one that wronged the claimholder. On the other hand, once a claim has in this way been flouted, it is no longer open to us to enforce or uphold it. There is no undoing the past, and if the wrongful action flouted a claim that was held against the agent, then the question of enforcing that very claim is now moot. The object of remedy can only be to make it, so far as it is in our power to do so, "as if" the violation had not

occurred, typically by compensating the wronged party for losses or harm suffered as a result.

My suggestion is that the secondary obligation to bring this state of affairs about should be understood as a residue of the original claim that was flouted. It is a residue of that claim in the specific sense that the relational structure of the secondary obligation derives from the relational structure of the primary wrong that gives rise to it, in particular from the fact that the primary wrong flouted a claim of the party to whom remedy is now owed.[45] What's done is done, and it is no longer open to anyone to honor or enforce the specific claim that was violated through the original wrongful action. But because that claim was flouted, the wrongdoer now owes it to the bearer of the claim to acknowledge wrongdoing, apologize, and (so far as possible) make amends.

A final point to emphasize about our accountability practices is that they have a constructive function that comes into clear focus when we appreciate the relational content of the moral norms that structure them. When reactive blame is considered in isolation, it can easily appear to be a punitive response that we would perhaps be better off without. But it is in fact one element within an ordered temporal sequence that gains significance from the relational norms that link it to other elements in the sequence. The immediate context of blame is the normative nexus between agents and the parties who have moral claims against them. The action that occasions blame is one that involves disregard for such claims, which to that extent reflects a deficient way of relating to the claimholder (a failure of what I have called interpersonal recognition).[46] In particular, it involves a failure to acknowledge the significance of the other parties' claims against one in one's deliberations about what to do, which may reflect a deeper failure to acknowledge the other parties as individuals with moral standing in the first place. Treating people in this way ruptures one's interpersonal relations with them, and this aspect of wrongdoing structures and renders intelligible the further reactions that it sets in motion. Reactive blame is rendered fitting by a failure of interpersonal recognition of this kind, and it in turn calls out for distinctive responses on the part of the wrongdoer, generating directed obligations of acknowledgment, apology, and repair that potentially culminate in the bestowal of forgiveness.

Seen in this light, blame appears to be a focused response to a rupture that has affected one's relationship to another party. It is occasioned by a failure on the agent's part to relate to the claimholder on a basis of interpersonal recognition. And it is embedded within a network of secondary obligations and entitlements that serve to repair the rupture and to reorient the parties toward

the norms that properly govern their interactions with each other.[47] The result, if all goes well, is to restore the value of interpersonal recognition that was damaged through the original act of wrongdoing.

3.4. The Relational Content of Blame

In the final section of this chapter, I would like to discuss the content of the reactions through which we characteristically hold people accountable for their actions. I have already noted that we customarily understand blame to incorporate an understanding of the relational requirements that structure our practices of interpersonal accountability. The aim of the present section will be to develop this point in greater detail, and to reflect on its implications for the relational account of morality.

It will be helpful to begin by contrasting the relational interpretation with a different approach that attributes to moral requirements a significantly social dimension. This is what I earlier called the social command theory, a version of voluntarism that explains moral obligations by appeal to the commands that are actually laid on agents by the members of the societies in which they live.[48] As we saw in chapter 2, voluntarist views in general connect obligations to the relationship that obtains between the agent and another party. But the relationship that matters is not that between a claimholder and a person against whom claims are held, but the relationship that is created when an authority addresses directives to a subject.

Like all versions of voluntarism, the social command theory grounds moral obligations in an antecedent interaction between different parties; in this case, it is the interaction through which commands are addressed by the members of a single society to each other. I argued in section 2.2 above that this kind of approach will be too crude to leave us with a plausible account of error with respect to moral obligation. From the fact that the members of a given community fail to exert pressure on each other to comply with a candidate requirement, it does not follow that the requirement does not represent a genuine moral obligation. Even if we abstract from this difficulty, however, I now want to argue that the social command theory distorts rather than illuminates the normative features of the obligations that are at issue.

Thus, if obligations are created by the commands of another party, their normative significance for the agent who is subject to them should derive from the relationship the subject stands in to the authority who issues the command. But this seems false to our intuitive understanding of the reasons at the

heart of morality. As an agent, my obligations to fulfill my promises and to respect the bodily integrity of other persons seem to derive from the interests of those who stand to be affected in one way or another by the things that I might do—from the promisee's interest in assurance, for example, or the interest of a potential victim in freedom from pain. This is connected to the point that honoring moral obligations involves a form of interpersonal recognition, through which one acknowledges the standing of other individuals as sources of claims. But the social command theory shifts the agent's focus in a different direction, toward the authority who issues the command in the first place. This seems to give agents the wrong kind of reason for complying with the requirements of morality.

If social commands take the form of exertions of social pressure (through blame and other expressions of opprobrium), agents will naturally think that their ultimate reason for complying with obligations is that doing so is necessary to avoid these sanctions.[49] True, voluntarists typically claim that commands generate obligations only when there is independent reason for agents to do the things that are commanded.[50] So the social command theory needn't deny that there are reasons for acting in conformity with moral obligations that are more directly connected to the interests of those who might be affected by the agent's behavior. Still, their status as obligations derives from the fact that directives have been addressed to the agent, through (for instance) the generation of social pressure by other members of the moral community. It follows that their distinctive normative significance, as obligations, is a function of the subject's relation to those who exert social pressure, which distorts rather than illuminates their reason-giving force. Voluntarists in this way substitute for the relation between agents and those who have claims against them the relation between agents and those who address commands to them.[51]

Social command theories also offer a distorted account of the reasons to which authorities are responding when they address demands to the agents who are thereby subject to them. Insofar as it is the actual addressing of the demand that creates the obligation in the first place, we cannot say that the authority's reason for addressing the demand is that it captures something the agent is morally obligated to do. Considerations of this kind are, as we saw in chapter 2, transparent to the reflections of the authorities whose commands give rise to obligations, on the general voluntarist approach. But the fact that the agent is under a moral obligation is, as it seems to me, the most natural reason for holding agents to a social demand that they should act in accordance with it. Nor is this problem avoided by those variants of the social com-

mand theory that appeal not to the actual demands people address to each other, but to the demands that it would be appropriate for them to address to each other.[52] The right kind of reason for addressing a demand of the relevant kind to someone, through the application of social pressure, is that the addressee is under a moral obligation to comply with it. But this consideration isn't yet available at this point in the voluntarist story, insofar as obligations are created by the existence of other kinds of reasons for addressing demands toward the agents whose actions they regulate.[53]

These problems are entirely avoided by the relational account of morality that I have been developing. Its notion of obligation is that of a complex of directed duties and claims, and these are considerations that are available to enter into the deliberations of both agents and the persons who have claims against them. As I noted in section 2.4, the existence of a nexus of this kind is naturally taken to have normative significance for both of the parties who are implicated in it. It enters into the reflections of the agent as a presumptive constraint on behavior, one that is connected in the right way to the interests of the people potentially affected by that behavior. A relational obligation is essentially one that is directed to another individual, who has a claim against the agent to performance. So in understanding themselves to be under moral obligations, on this way of conceptualizing them, agents are thinking of the implications of their actions for the interests of other individuals, and acknowledging the standing of other individuals as sources of claims against them. By contrast with the voluntarist approach, this traces the normative significance of moral obligations to the right source, and connects compliance with such obligations to an attractive positive ideal of interpersonal recognition.

The relational interpretation also gives the notion of moral obligation work to do in the perspective of the members of a society who hold each other morally to account. In adopting this stance toward each other, people are exerting a kind of social pressure that helps to constitute them as members of a common moral community. Our reasons for participating in this interpersonal practice, however, derive from the independent fact that there are complexes of directed obligations and claims in place that antecedently connect us to others in a normative nexus. We understand that there are things that others owe it to us to do, and this makes it reasonable for us to take a special interest in whether they live up to the relational standards that specify these obligations, holding them accountable with respect to the obligations. Similarly, disregard for these relational standards provides a normative basis for reactive and other forms of blame, insofar as it amounts to disregard for claims we have

against agents to compliance with those standards. It is because there are things that people owe it to us to do that we hold them to account when they fail to live up to their obligations and claims. The reasons to which our accountability practices are responsive are therefore reasons of the right kind, stemming from facts about the moral obligations and claims that link people to each other in a normative nexus.

These remarks help to fill in the picture that was drawn very swiftly at the end of chapter 2, of moral obligations as considerations that have normative significance for two different parties. In conclusion, I now want to look more closely at the attitudes on the part of claimholders that are rendered appropriate by the obtaining of directed moral obligations. I have characterized these attitudes so far in terms of interpersonal accountability. They involve a stance of holding agents to moral demands in a way that confers on them responsibility with respect to the demands, and that is connected with distinctive responses of moral blame on occasions when the demands are breached. The general idea, then, is that directed obligations make it fitting for claimholders in particular to adopt this array of characteristically accountability-conferring stances toward the agents who are under moral obligations.

If this is right, however, then we should expect our interpersonal accountability practices to prioritize the relationship between agents who are subject to moral requirements and the claimholders to whom those obligations are owed. Claimholders, after all, are in a privileged position by comparison with the other people who might be affected by the agent's actions. If A flouts a duty that is owed to B, then it is B in particular who is wronged or has suffered a moral injury, and B would seem to have a specific grievance or complaint about what A has done that is not shared with other parties. This is connected to the idea that it is B, in the first instance, who has an immediate reason to respond to A's action with blame. Our accountability practices should reflect these facts about the claimholder's privileged normative position in the wake of wrongdoing; or at least this is how they should look if I am right in suggesting that they are structured through relational moral norms.

One might express the point in terms of an analogy between moral responsibility and the law. There are two legal paradigms to which we could appeal in thinking about our interpersonal practices of moral accountability. One of these is the criminal law, where agents of the state are authorized, on behalf of the community for which they work, to bring charges against defendants for violations of relevant statutes. The other paradigm is private law, where standing is assigned to individual claimholders who may sue for damages, on their

own behalf, that result from actions that wrong them in particular. If we think in terms of the criminal law model, then the paradigmatic expression of moral blame should be a stance that is taken up on behalf of the larger community of moral persons, and that it is in principle open to anybody to adopt. To blame another, on this conception, is to assert moral demands impersonally as a representative person, acting out of the interest that we all have that relations between us should be conducted on a moral basis.[54] On the private law model, by contrast, the paradigmatic expression of blame is the interpersonal assertion of claims that one holds in one's own person in cases where those claims have been disregarded. In these terms, the hypothesis expressed earlier would be that the private law model should more closely reflect our moral practices of accountability if I am right to contend that they are structured through essentially relational requirements.

But this hypothesis seems consistent with a plausible independent analysis of our accountability practices. Thus, the reactive attitudes that are widely taken to be paradigmatic forms of moral blame in fact have a content that reflects the privileged position of claimholders. The primary reactive attitude is perhaps resentment, and as noted in section 3.2, this is an emotion that it is open to people to feel only on occasions when they believe themselves to have been wronged in some way. You can't really resent the morally problematic thing someone has done, however reprehensible it might be, if you don't understand the action to be one that violated some specific claim that you had against the agent not to do it. Resentment is thus the accountability-ascribing attitude that is rendered distinctively appropriate on the part of claimholders when agents flout their claims. It precisely reflects a claimholder's privileged complaint about what the agent has done, in a way that mirrors the positional conception of standing within private law.

To privilege resentment in this way in an account of interpersonal accountability, however, is not to hold that blame is available only to the party whose claims have been flouted. What we should expect, rather, is that other forms of blame, such as those available to third parties or to agents who hold themselves to account, are parasitic on our understanding of the reasons that those who are wronged have for resenting what was done to them. But this corresponds well to common ways of conceptualizing the more extended forms of reactive blame. Thus indignation is plausibly characterized by P. F. Strawson as the "vicarious analogue" of resentment; it is an attitude that we typically feel on behalf of the person who has been wronged, and it in this way seems to presuppose that the action we are objecting to violated a claim that was held

by another party.[55] The angry disapprobation we experience in a case of this kind rests on the attribution to the claimholder of a privileged basis for complaint against the agent, one that would render resentment fitting on the claimholder's part. Third parties do not always have a good reason for reactive blame in response to wrongdoing; but when they do, it is natural to understand their indignation as derivative from the complaints that wrongful action generates in the person whom it wrongs.

On the criminal law model, by contrast, we would expect an impersonal form of indignation to represent the primary expression of moral blame. An impersonal conception of blame, in the relevant sense, would be based in interests we all share in the maintenance of a well-ordered moral community, rather than in the claims of individuals that might provide a privileged basis for them to object to wrongdoing on their own behalf. As we saw earlier in this section, however, an impersonal conception of this kind has difficulty understanding the reactions that it defines as responses to the flouting of obligations that are antecedently intelligible. The natural development of this approach is a kind of social voluntarism that sees the assertion of moral demands, by the representative moral person, as the original source of moral obligations, rather than a response to their violation. On the private law model, by contrast, the directed duties that give claimholders a privileged basis for complaint also define a notion of obligation that is available to figure in the thoughts characteristic of third-party blame. Indignation is about the flouting of obligations, precisely insofar as it is a vicarious reaction that is parasitic on the standing of claimholders to object on their own behalf to what was done to them.

It is a further reflection of the applicability of the private law model that indignation does not go together with the special authority to forswear blame that we have seen to be characteristic of resentment. Special contexts aside, third parties do not really have the normative standing to forgive the wrongdoer whom they might blame; this is among the positional features of forgiveness that, as I have argued, presuppose the relational approach to the requirements that structure our accountability practices. But it is unclear why there should be this asymmetry in normative power if third-party indignation, on behalf of the moral community, is the paradigmatic expression of blame. Within the criminal law model, there is room for those who blame on behalf of the moral community to overcome their angry reactions. But this will have the character of mercy or compassion rather than forgiveness, which is avail-

able only to a person who is understood to have a privileged basis for complaint about what the wrongdoer did.[56] The relational interpretation of indignation, as an essentially vicarious attitude that is adopted on behalf of someone who has the primary standing to blame, better accommodates these aspects of interpersonal accountability.

Consider next self-blame, the canonical expression of which is guilt. One salient feature of this reaction is that it reflects an understanding of oneself as not only deficient, relative to an ideal of personal attainment, but as having violated obligations or practical requirements that are binding on one.[57] As we saw in the preceding chapter, however, the relational interpretation of moral rightness provides a natural account of its standing as a source of practical requirements of this kind. Our sense of ourselves as having flouted requirements in these cases might well be understood to reflect our recognition that we owed it to another party to comply with them. In the same spirit, guilt is often said to involve the awareness on the part of agents who have done something wrong that they have caused a rupture or break in a valuable relationship that they stand in to another party.[58] This seems to correspond intuitively to the idea that the moral requirements the guilty agent is understood to violate are ones that are owed to another individual in particular. The rupture one has brought about, we might say, is precisely the fact that one has acted in a way that gives the other party a privileged complaint, something that would make it fitting for that party to resent one for what one has done.

Guilt and indignation, on this plausible way of thinking about them, incorporate an implicitly relational understanding of the moral requirements that help to fix their contents. Once we appreciate this point, however, then the possibility opens up of extending reactive blame even further, to cases involving parties who may not be in a position to assert claims or to understand when they have been flouted in their own case. It is natural to suppose that young infants, for instance, or adults who are severely mentally impaired, have claims just like the rest of us do to be treated with consideration and due regard—this despite the fact that they may not be in a position to hold other people to the demand that their claims be acknowledged and respected, or even to understand very clearly when these standards have not been met. Although it is a more controversial position, one might finally want to broaden the class of moral claimholders further to include nonrational animals, at least if they have a developed capacity for pain and suffering and reasonably complex interests.[59]

Claims, we have been assuming, have normative significance for the claim-holder; but in cases of all these kinds, claimholders are not actually in a position to adjust their attitudes in response to the specific recognition that reasons for such an adjustment are in place. Here, it seems to me, the notion of a vicarious reactive attitude has a further potential application. A third party, acting as it were as a trustee for the claimholder, can assert a claim on behalf of the injured party through attitudes such as indignation, which incorporate an understanding of the action that is blamed as wronging the injured claim-holder.[60] In doing this, we are incorporating the individuals in question into our practices of interpersonal accountability, insofar as we are asserting claims on their behalf that they are not in a position to assert themselves. This is a perfectly coherent move, it seems to me, involving a distinctive deployment of the capacity for vicarious blame that is already available to third parties in cases in which the individuals who were wronged are able, in principle, to assert claims for themselves.

To this point, I have mainly focused on the reactive forms of blame, arguing that the paradigmatic instances of angry disapprobation are naturally understood to involve an implicitly relational conception of the wrongs to which they are a response. It is worth adding, however, that the paradigmatic nonreactive forms of blame similarly seem to rest on a relational conception of those wrongs. T. M. Scanlon, for instance, has influentially argued that blame should primarily be understood to involve the adjustment of one's intentions in response to reasons for such adjustments that are provided by the attitudes of another party.[61] Scanlon argues, more specifically, that attitudes of disregard for moral standards typically impair the agent's relations to other persons, and that the impairment of a relationship gives the other parties to it reasons to modify their behavior in response, including reasons to avoid the agent, to withdraw from trust-based interactions with the agent, and so forth. These are reasons that one might act on even in the absence of the attitudes of angry disapprobation characteristic of reactive blame.

I think Scanlon's talk of the impairment of relationships in this context is potentially misleading, since it suggests that nonreactive blame is available only to parties who already stand in an ongoing historical relationship with each other, one that is eligible to be impaired by the attitudes of the parties to it.[62] Abstracting from this issue, however, I am prepared to acknowledge that there is a characteristic syndrome of behavioral responses to wrongdoing that might be understood as an extended form of blame. Note, however, that these

nonreactive instances of blame make sense only insofar as they are responses to the flouting of essentially relational requirements.

We have reason to withdraw from interactions with other people who act wrongly, because their doing so reflects attitudes of disregard for us and for the claims we have against them. Such attitudes may not impair a relationship between us that exists antecedently; but they are certainly attitudes that have personal significance for us, on account of the fact that they are attitudes of targeted disregard. The syndrome of withdrawal, retrenchment, and avoidance behaviors is a warranted response to actions that reflect negative qualities of will toward us. As I have argued, however, the flouting or neglect of moral requirements will have this character only if we understand those requirements in relational terms. It simply doesn't make sense to respond to wrongdoing with nonreactive blame except insofar as it expresses attitudes of ill will toward us in particular, in ways that rest on a relational understanding of the norms against which the wrongdoer has offended.[63] This is a further reason for thinking that our accountability practices should be understood in terms of the private law model sketched above.

In his classic paper "The Nature and Value of Rights," Joel Feinberg invited us to think about "Nowheresville," a hypothetical community whose members possess all of the familiar moral notions except that of a right.[64] In discussing the elements that would need to be added to introduce rights or claims into this community, Feinberg identifies as crucial a "performative" sense of asserting a claim, which he understands as a "rule-governed activity" that is "public, familiar, and open to our observation."[65] His ultimate suggestion is that rights just are considerations that can validly be put forward as claims in this performative sense, and that they acquire their meaning and interpersonal significance from their relation to this practice. We can understand this as an argument that starts from the notion of a moral right or claim, and works backward to identify the forms of social interaction that are necessary to make sense of these normative notions.

In this chapter, I have mounted a complementary argument that locates relational moral notions within the distinctive practice of interpersonal accountability, understood as a practice that is organized around reactive and other forms of blame. My starting suggestion, put forward in section 3.1, was that it is a constraint on a nondebunking interpretation of morality that it should make sense of the idea that moral norms are distinctively suited to structure interpersonal accountability relations of this kind, so that disregard

for them would provide a normative basis for the various manifestations of blame. The first phase of my argument, completed in section 3.2, established that the relational interpretation of morality is tailor-made to meet this constraint, insofar as it connects moral obligations constitutively to the claims that individuals have against them. It in this way does justice to our understanding that morality is an inherently social phenomenon.

In the argument's second phase, I showed that our interpersonal practice of holding people accountable for complying with moral requirements itself presupposes a relational understanding of them, as obligations that are owed to individuals who have claims against us to their compliance. This idea was developed in sections 3.3 and 3.4, which explored the positional character of the interlocking elements that together make up our practice of accountability, as well as the relational content of reactive and other forms of blame. Seen in this light, the practice of accountability may be understood to involve the assertion of essentially interpersonal claims (to echo Feinberg's language), together with focused responses on the part of wrongdoers to such assertions. If the relational account is rejected, then we shall have to conclude that morality is not everything that it seems to be, insofar as standards of rightness will not provide a normative basis for relations of interpersonal accountability; and those practices, for their part, will similarly be called into question, insofar as their content and structure show them to be organized around inherently relational norms.