

Chapter 1

Why Be Rational?

Does man think, then, because he has found that thinking pays—because he thinks it advantageous to think?

Ludwig Wittgenstein (1958, p. 134)

I. A Source of Normative Force

Suppose a student comes to my office asking which courses she should take next semester. I respond by asking her what she wants to do. She says she wants to major in philosophy. I ask why. She answers in a way that reassures me that she has no illusions about the philosophy major. She has a sense of what it is like to study philosophy and understands that she has only a slim chance of being able to make a living as a professional philosopher. Knowing all this, she definitely wants to be a philosophy major.

So, I tell her she ought to meet the logic requirement. Now I take it that what I told her is true. I could be wrong about the logic requirement, but if I am wrong, it is because I have made a factual error. (Perhaps my department does not have a logic requirement.) Although I might be wrong about what she ought to do, the fact remains that, because she has ends, there are things she ought to do. Given that she truly has the goal of majoring in philosophy and that meeting the logic requirement truly is part of what she has to do to complete the major, she truly *ought* to meet the logic requirement.

Something in the situation allowed us to derive a normative conclusion, a conclusion about what she ought to do. We did not formally derive an ought from an is. Nor would there be much point in trying. As the student and I sit in my office pondering her future, the so-called is-ought problemⁱ is the least of our concerns. The is-ought problem is a philosophical analog of a real problem, and the two of us are pondering the

real thing, not the philosophical analog. The real problem involves deciding what she should do given the information at hand. But although we may not have derived an ought from an is, we did derive an ought.

What allowed us to conclude that she ought to meet the logic requirement? Two things, mainly. First, we had information about her ends. Second, we had information to the effect that meeting the logic requirement would be a means to her ends. We derived a normative conclusion from facts about a person's ends and facts about what would achieve those ends, which means that facts about ends have a certain potency, a certain *normative force*: they give us reasons for action.ⁱⁱ In particular, to have goals is to have reasons for acting in one way rather than another. If acting in one way rather than another would serve my ends, then I have reason to act in one way rather than another.ⁱⁱⁱ

This is one source of reasons for action: if doing X would best serve a person's ends, then the person has a reason to do X. If this is not the only source of reasons for action, then we may need to qualify the last statement, adding the phrase "all other things equal" in recognition of the possibility that other normative considerations could conflict with means-end considerations. Whether or not other sources of normative force exist, though, we have identified at least one way in which normative considerations come into existence. They come into existence along with ends.

Are ends the only things that give us reasons for action? Or can we have reasons for action independently of whether the action would serve our ends? Part of this book is about how other kinds of normative force—including the peculiar kind of ought associated with moral constraints—could emerge from such humble beginnings. Suppose we begin with a rudimentary self-regard borne of biological instinct, and an ability to evaluate alternative means to our ends. Beginning from there, how do we come to have other kinds of reason for action?

How do people come to care about themselves in the way we care about ourselves? This may seem like something to take for granted, but the easiest things to take for granted

are sometimes the hardest things to explain. It is easy to explain in biological terms why we want to survive, or to have sex, but some of us want people to spell our names correctly. Some of us want dishes that match. Why?

How do people come to care about others in the way we care about others? Again, it is easy to explain in biological terms why we care about the survival of our children, but some of us care who wins the World Series. Our concerns sometimes extend not only beyond our children, but even beyond our species. We feel sickened by the extinction of nonhuman species even when they have no commercial value and are found only in countries we will never see. Why?

I said our goals give us reasons for action. I say this without presuming that the goals themselves are reasonable. In fact, later chapters describe a process by which our goals can become subject to rational critique—and thus how we acquire reasons to have some goals rather than others. There may even be ends, like the end of completing a philosophy degree, that we do not acquire until we become convinced that pursuing such an end is reasonable. Nevertheless, once we have an end, simply having it gives us reasons for action. For example, whether or not we can justify having a goal of survival, it remains that if our goal is to survive, there are certain things we have reason to do. Accordingly, we can make sense of the fact that even biologically given ends, like the end of survival, imply reasons for acting in one way rather than another. To have an end is to have something with normative force, something that gives us reasons for action, regardless of whether we in turn have reason to have the end.

Of course, not all reasons for action are equally powerful. Further, questioning our ends may sometimes serve our ends, and if the exercise reveals that we have no good reason to be pursuing a certain end, that can undercut the end as a source of reasons for action. The following chapters have more to say about the rational critique of ends. This chapter discusses rationality as a particular kind of normative consideration.

II. Endorsement

What do we mean when we call something rational? My answer is distinct from, yet inspired by, Allan Gibbard's. Gibbard (1991, p. 6) says that to call something rational is to endorse it. More precisely, according to Gibbard, to call something rational is to express acceptance of norms that permit it (p. 7).

One attraction of Gibbard's approach is that it directs attention away from the thankless task of defining the term 'rational' as it appears in ordinary language. Gibbard invites us instead to focus on the connection between being rational and having reasons for action. Some of the details of Gibbard's approach are worth challenging, though. For example, when Kate calls X rational, is Kate really expressing *acceptance* of norms that permit X, or is she merely expressing a belief that X is permitted by certain norms? More generally, must I accept a norm in order to express my belief that it permits X? I think not. If my dinner companion seems flustered about conventional table manners, it would not be strange for me to say, "That's the right fork, for what it's worth." In this way, I genuinely endorse her choice of forks by expressing my belief that her choice is permitted by norms of etiquette that *she* accepts (and about which she would like to be more knowledgeable). But I express my belief without implying that I accept those norms myself. I merely know about them.

Thus, in matters of etiquette, using a norm as a basis for endorsement does not presuppose that one accepts the norm. The same holds true when we endorse a choice as rational. If my dinner companion tells me she has accepted a proposal of marriage, it would not be too strange for me to convey, if not in these exact words, "I'm happy for you. For you, it's rational, because you accept legal and religious norms that constitute the institution of marriage. I reject those norms myself and thus will never marry, but I can see that for you getting married makes perfect sense." So, although Gibbard's analysis in terms of norm acceptance is not one I reject out of hand, I observe that people can call something rational

with more detachment than his analysis seems to allow. Accordingly, I will stick with the more general of Gibbard's two formulations: to call something rational is to endorse it.

Although I prefer the more general formulation, it is in one respect too general. Gibbard is right that calling X rational falls into the genus of ways of endorsing X. But it is less clear that calling X rational is coextensive with that genus. When we call something rational, are we merely endorsing it in a generic way, or are we offering a specific kind of endorsement? The latter seems more natural. At a restaurant, when the waiter brings the salad and my companion hesitantly picks up the smaller fork, I might say, "That's the right fork." Am I calling it rational to pick the smaller fork? Not to my knowledge. The norm by which I endorse her selection of the smaller fork is a norm of *etiquette* rather than of rationality. I can express a norm of etiquette, one that mandates using the small fork to eat salad, without expressing or implying the thought that using the small fork is rational. I can endorse X (and even express acceptance of norms that permit X) without at the same time calling X rational.

I also suspect that our calling something rational does not quite entail that we endorse it. As Gibbard notes, we can call something rational in a sour tone of voice (p. 50). On the other hand, we can offer an endorsement without being cheerful. A grudging endorsement is still an endorsement. Perhaps the thing to say is that the endorsement we express when we call a choice rational need not be unconditional or without reservation. In part, this depends on whether we take the term to express an "all things considered" judgment or something more narrow. If in calling a choice rational we mean that it is the thing to choose all things considered, then we can have no reservations about it beyond the reservation that we might be mistaken about it being rational. In contrast, if we express a less comprehensive kind of endorsement when we call a choice rational, then we may have reservations about it (and may convey that fact by speaking in a sour tone of voice) despite being convinced that it is rational.

I do not want to leave the impression, however, that something momentous hangs on how we decide to use the term ‘rational.’ Gibbard says that settling what it is rational to do settles what to do, period (p. 49). I think we could go either way. For my purposes, it makes no difference, because regardless of whether we follow Gibbard in reserving the term ‘rational’ to refer to “all things considered” judgment, it remains the case that moral considerations sometimes conflict with other considerations when we are deciding what to do. Such conflicts are internal to rationality in Gibbard’s all things considered sense. But they still occur. Regardless of whether we identify rationality with the entire genus of normative endorsement, we are left needing to weigh various considerations, including moral considerations, when deciding what to do. Section VI says more about how moral endorsement differs from other kinds of endorsement. The next two sections discuss when something warrants the endorsement we imply when we call it rational.

III. Warranted Endorsement

Gibbard says that when we call something rational, we are doing more than offering a mere description of it. His insight here is crucial. Calling something rational is also (or at very least tends to be) an expression of endorsement. Gibbard chastises people who “take the characteristics that go to make something rational and have them constitute the meaning of the term” (p. 10). In my words, Gibbard’s point is that the descriptive features that occasion endorsement are less central to the meaning of rationality than is the element of endorsement itself. People quarrel over whether particular choices are rational, and generally the dispute is not terminological. If my student argues with her parents about whether it is rational for her to major in philosophy, they are not arguing about what the term means. They all know that to call a choice rational is to endorse it. Their argument is about whether majoring in philosophy is worthy of the endorsement that they all know is implicit in calling it rational.

This is plausible, as far as it goes. Yet, to call a choice rational is to try to say *something* about it. Calling a choice rational is not mere description, but neither is it mere endorsement. I said that Gibbard views the term ‘rational’ as not merely descriptive. Actually, Gibbard thinks the term is not descriptive at all. Gibbard insists that to call a choice rational is not to attribute a property to it, not even the property of being permitted by accepted norms (p. 8). Now, if he were saying there is no particular property in terms of which we can analyze the meaning of the term ‘rational’, he might be right. But the real question concerns whether there are qualities in a choice to which we *refer* when we pick it out as rational. And if not—if there is nothing in the choice that leads us to call it rational—then there is no reason for our endorsement. In that case, something has gone awry.

The point for us to stress, then, is that calling X rational does more than simply convey endorsement. It also conveys that we think endorsement is warranted. Even if Gibbard is correct in saying that no particular property ascription is implied by the term’s meaning, we nevertheless are not using the term properly unless we take X to have some property or other that gives us a basis for endorsement. Think about it from a listener’s point of view: when we call something rational, people understand us to be saying not just that we endorse it but also that we think we have *reasons* for endorsement. To call X rational is not to endorse X in the way one might whimsically endorse butterscotch as a flavor of ice cream. Rather, to call X rational is to say X warrants endorsement. Listeners may not interpret us as having *specified* our reason for endorsement, but they will interpret us as thinking we *have* a reason. So I accept Gibbard’s claim that to call X rational is to endorse it, but that leaves open what, descriptively speaking, our endorsement is supposed to be tracking when we call X rational.

Here is one answer. Sometimes, at least, we call a choice rational because we think it will serve the chooser’s ends. Understood in this way, to call a choice rational is, first, to endorse it, second, to have a reason for endorsement, and third, to have as one’s reason for endorsement that the choice will serve the chooser’s ends.^{iv} When I call a choice rational, I

may be saying that it will in fact serve the chooser's ends or that the chooser has good reason to think it will (whether or not it actually does so). In any case, rational choice, so understood, involves seeking to choose effective means to one's ends. This is the heart of my characterization of rational choice. It does not have a lot of content, but we will see that it has enough to support the conclusion that, for most people, being moral is rational.

IV. Beliefs as Inputs and Beliefs as Outputs

Whether the chooser's *beliefs* are rational is another issue. What an agent has reason to choose depends on what the agent believes, but I say this without presupposing any particular theory of rational belief. Practical rationality, or rational choice, is about making choices in light of what one believes. Being rational in the practical sense involves seeking to choose effective means to one's ends, given one's beliefs. I acknowledge that issues of epistemic rationality arise at an earlier stage when one forms beliefs. If we were discussing belief rather than choice, our grounds for endorsement might well be something other than means-end efficacy.^v

However, although a theory of rational choice need not be committed to any particular theory about rational belief, I do want to comment on a curious property of beliefs that is crucial to any theory of rational choice aspiring to have practical significance for human beings. What I have in mind is that the processes by which we acquire and revise beliefs can be compromised by our choices. For example, when we make choices we could not justify to others (or would feel uncomfortable trying to justify), such a state tends to induce irrational belief-forming processes: projection, repression, spurious rationalization, and so on. Clearly, such processes can hurt us. They can obscure the nature of our ends. They can lead us to overlook incongruity in the means we choose to our ends. They can block our pursuit of self-understanding.^{vi}

Because the consequences of irrational belief-forming processes can be momentous, it is of the essence of humanly rational choice that it involves recognizing and coping with the fact that we cannot take our own rationality for granted. If we care about rational choice, we have to care about the rationality of our belief-forming processes to the extent that their reliability is predictably affected by our choices. In a discussion of moral motivation, T. M. Scanlon observes that the desire to be able to justify one's actions to others on grounds they could not reasonably reject may not be universal, but it is quite strong in most people (1982, p. 117). I mention the point here because, when we feel others *could* reasonably object to a choice, that feeling can count against the choice's rationality, insofar as such feelings tend to trigger psychological defense mechanisms that compromise the reliability of our belief-forming processes.^{vii}

In different words, practical rationality proceeds from whatever beliefs a person has, and so the epistemic rationality of beliefs as inputs into choices is a peripheral issue here. However, we do need to distinguish between beliefs (as inputs into decisions) and changes in beliefs (as outputs of decisions). The epistemic rationality of our beliefs as evolving *outputs* of our choices is a central concern practical rationality.

For example, when a choice induces spurious rationalization as part of its output, it thereby compromises the reliability of the person's beliefs as inputs into future choices; it lowers the quality of belief-inputs into future choices, which typically would be contrary to practical rationality. Therefore, even in the most narrowly practical terms, we cannot ignore issues of epistemic rationality. There is no need here for a *theory* of epistemic rationality, but we do need to be aware that, human psychology being what it is, our choices affect the reliability of our belief-forming processes. This is just one of the many ways in which our choices affect us, though, and that more general topic—our choices and how they affect us—is this book's main concern.

If we needed to formulate a sufficient condition for rational choice, we might need to say more about the role of beliefs as inputs into our choices. We might need to say

something like: a choice is rational if it serves an agent's ends, assuming the agent is cognizant of how it will serve his or her ends and makes the choice on that basis. Fleshing out such a qualification might call for a theory of epistemic rationality. But no matter what caveats we might need to introduce if we were attempting to formulate a sufficient condition, there remains a close link between rational choice and means-end efficacy.

We can describe that link by saying means-end efficacy is a *supporting condition* for rational choice.^{viii} A supporting condition is a qualified sufficient condition, qualified in the sense of being sufficient in the absence of countervailing conditions. In different words, a supporting condition is sufficient to shift the burden of proof. Analogously, in a court of law, certain kinds of evidence are regarded as sufficient, in the absence of countervailing evidence, to establish legal liability. For example, under a certain set of liability rules, if the plaintiff can show that you ran over his bicycle, that suffices to establish your legal liability unless you can show in turn that the plaintiff was negligent in leaving his bicycle on the ground behind your car. In the same way, evidence that a choice will serve the chooser's ends warrants our endorsing the choice as rational, barring countervailing evidence (for example, barring evidence that the chooser was not cognizant of the manner in which the choice would serve his or her or its ends).

In this qualified way, the fact that a choice will serve the chooser's ends suffices as grounds for endorsing the choice as rational. Of course, actual means-end efficacy is neither necessary nor strictly sufficient as grounds for calling a choice rational, but we do not need necessary and sufficient conditions.^{ix} When we see people choosing effective means to their ends, we have grounds for calling their choices rational, so long as we recognize the possibility of countervailing evidence. For example, we might resist calling a choice rational, even though it served the chooser's ends, if the chooser had no idea that it would serve his or her ends in the way it did. An objectively warranted choice may yet have been unwarranted in a subjective sense.

V. Objectively Warranted Endorsement

Previous sections sketched the view that to call a choice rational is to endorse it as a means to the chooser's ends. This section distinguishes between warranting endorsement in an objective sense and warranting endorsement in a subjective sense, then considers how the distinction leaves a strategy of maximizing expected utility open to rational criticism.

Perhaps endorsement can be warranted only from the perspectives of particular people. I judge my choices rational according to whether they warrant endorsement from my perspective. I judge your choices rational according to whether I think they warrant endorsement from your perspective. Be that as it may, the point remains that to endorse a choice as rational is to endorse it not arbitrarily but on particular grounds. If we endorse a choice as serving the chooser's ends, then whether the choice actually warrants such endorsement is a matter of fact, not a matter of opinion. This is not to deny that this matter of fact might turn on facts that are agent-relative. Whether choosing to pursue a career in philosophy will serve my student's ends may be a factual question, but the answer will of course depend on facts about her. The rationality of her choice does not turn on her opinion *per se*, and thus is to that extent objective, but it does turn on her ends. Thus, given my means-end conception of rationality—a conception that can hardly be considered unusual—it follows that whether a choice warrants endorsement as rational is a question the answer to which is both agent-relative and objective.

Why does this matter? Gibbard says we do not need objective standards in order to get guidance. Consider a variation of an example of Gibbard's (p. 18). A hiker is lost in the woods, wanting very much to get back to town. There is an objectively best path in this case, for there is a path that as a matter of fact would best serve the hiker's ends: a straight line back to town. Nevertheless, the rational path in this case is not identical to the objectively best path. The rational path is the path that follows the river. Why? Because the hiker is operating under incomplete information. (He is, after all, lost.) Because he has

no idea which direction constitutes a straight line back to town, he concludes that following the river has the highest expected utility. Although the river path is a roundabout way of getting back to town, it is the path most likely to eventually bring him safely back. The river path is thus subjectively justified. It is subjectively justified in the sense that, given what the hiker knows about the situation, the river path seems most likely to serve his ends. If it actually serves his ends, then it is an objective success as well.

What does Gibbard conclude? Gibbard reasons that either we do not know enough to follow an objectively best path, one that best serves our ends, or if we do know enough, then the path we follow is also subjectively best. Either way, what we follow is the subjectively best path. Objective standards are either useless or superfluous (p. 43).^x Gibbard seems to think nothing more needs to be said, that the idea of a subjectively right path is intelligible independently of objective standards of rightness.

Not so. On the contrary, the idea that there is a subjectively best path presupposes that the hiker has standards of objective rightness. Suppose the story continues. The hiker's downstream trail leads him ever deeper into the woods. The sun goes down, his legs will carry him no farther, and it begins to look as though he will freeze to death. At this point, does the hiker review his strategy, find that it was impeccable by subjective standards, and die happy? Not at all. In fact, he has cause for regret, for his subjectively justified strategy is an objective failure. His strategy had the highest subjective probability of success. The point of the subjectively justified strategy, though, was to maximize the likelihood of objective success, where objective success consisted of getting safely back to town. A subjectively right strategy is a means to the end of getting the hiker back to town, and whether it serves that end is an objective matter.

Gibbard thinks there is "no need, then, for an account of what it means to call an act wrong in the objective sense" (p. 43). However, as I understand the term 'objective' (Gibbard must have something else in mind), this is not true. On the contrary, we need a sense of objective right and wrong to have a *purpose* in formulating subjective strategies,

and to distinguish between the successes and the failures of subjectively justified strategies. It is important to be subjectively right (i.e., to choose a strategy that, to the best of one's knowledge, is most likely to achieve the end) if and only if it is important to be objectively right (i.e., to actually achieve the end). Maximizing expected utility is a means to the end of maximizing actual utility. More generally, being subjectively justified is a means to the end of being objectively right.^{xi}

Two points of clarification. First, I am not arguing with Gibbard about whether the hiker is rational. I agree that the hiker is rational and that he is rational precisely because his strategy is subjectively justified (that is, it has the highest subjective probability of getting him back to town). I claim, though, that objective standards of rightness, far from being superfluous, are crucial to an understanding of subjective justification. There is a distinction, not between subjective and objective rationality as such, but rather between the subjective justification we can have in choosing means and the objective success we can have in achieving ends.^{xii} There can be no conception of a strategy's subjective justification except in the context of a conception of the strategy's objective success.

Second, saying this does not commit us to holding that some ends are better than others independently of people's preferences. Some ends may well be better in that sense, but I have not relied on any such idea here. Suppose we grant that what it is rational for a person to choose ultimately depends on the person's actual present purposes. Consider what this supposition does not imply. It does not imply that people, as rational agents, must take their present purposes as a simple given. On the contrary, it is a fact about human beings that our ends can change, which suggests that a person's existing corpus of ends sometimes might be well-served by processes that predictably cause the corpus of ends to change. For example, our desire to be healthy might be well-served by developing nutritional and exercise goals. Further, a preexisting end of being healthy is best served if we come to like jogging and eating broccoli for their own sake, not only for the sake of

their conduciveness to our health. Taking our present ends as simple givens could lead us to overlook opportunities to further them.

Even more puzzling, there may be cases in which one has no real means of weighing alternatives until after one makes some basic choices. Even when there is nothing prospectively to be said for one alternative compared to another, it can still be true that whichever one we choose, we will in retrospect be glad we chose it. This is because our choices can change us in such a way that we come to have a framework within which previous choices emerge as suitable for the people we are becoming. In simpler words, we tend to accommodate ourselves to our choices; we grow into them.^{xiii}

Static rationality involves assessing means to present ends we take as given. We can talk about rational choice in a static sense, so long as we realize that static rationality is not the only thing there is to talk about. There is a conversation that comes to an end when we have identified means to the chooser's actual present ends. But there is another conversation that begins when we reflect on how people's present capacities and present ends evolve as a consequence of their choices. This too is a conversation about rational choice, but it is rational choice of a different kind, a reflective and dynamic kind of rationality that we understand better if we keep in mind that it is different.

VI. "Why Be Rational?" As an Empirical Question

Because there is a difference between being subjectively justified and being objectively right, we cannot jump from saying rational choice involves seeking to choose effective means to one's ends to saying rational choice involves seeking to maximize expected utility. Consider: why would anyone seek to maximize expected utility? Presumably, one thinks that maximizing expected utility will have actual utility, which is to say, one thinks that maximizing expected utility will serve one's ends. Maximizing expected utility is a strategy. A strategy is something we adopt as a means to an end. Whether a

strategy serves its purpose is an objective and contingent matter. Therefore, there is no strategy, not even maximizing expected utility, that we can simply equate with being rational.^{xiv} Even if we can say the hiker's end is to maximize actual utility, it remains an open question whether a strategy of maximizing expected utility is an effective means to that end. Maximizing *expected* utility cannot possibly be the hiker's end; by hypothesis, he *did* maximize expected utility, but (in the version of the case where he became even more lost) he did not get what he wanted.

Part of the point of asking "why be rational" is to draw attention to the fact that it is tautologous that we have reason to do what serves our ends and yet it is not tautologous that we have reason to be utility maximizers or to affect a deliberate and calculating demeanor. To endorse a choice as rational is to endorse it as a means to an end. This is the basic concept. To get to the conclusion that it is rational to employ a utility-maximizing strategy or to affect an unemotional demeanor, we need empirical premises, premises that are sometimes true in the real world but that often are not. If rational choice is thought of as involving a certain way of pursuing goals, then it will be an empirical question whether people are better off pursuing goals in that way.

For example, there are limits to the rationality of being deliberate and calculating. Some situations call for spontaneity or for a certain lack of concern. A well-ordered life makes room for a certain amount of disorder, both because some disorder is inevitable and also because good things sometimes come from it.^{xv} Effectively serving one's ends need not always involve being fully in control of the situation or even being fully in control of oneself. Some of life's most precious moments are moments when we are merely "along for the ride." For any strategy or character trait having a close empirical association with rationality, there will be limits to how rational it is to cling to that strategy or that persona, for there will be circumstances in which things that normally serve our purposes would be counterproductive.

Being rational typically involves carefully weighing costs and benefits, but carefully weighing costs and benefits is itself an activity with costs and benefits, and so if we step back to weigh the costs and benefits of being rational, we may foresee situations in which being rational in that sense is not worth the price. There might be times when we will lose opportunities if we stop to carefully think them over. There may also be times when stopping to think them over would change the nature of the alternatives. For example, buying your spouse a bracelet only after carefully deliberating about anticipated costs and benefits might not be the same kind of act as buying your spouse a bracelet in a burst of carefree affection. The latter kind of act might serve your purposes better.

So much for things contingently associated with being rational, things like hard-headed strategies and calculating demeanors. These things usually serve us well, but are sometimes out of place. Setting aside contingent associations, then, rational choice essentially involves means-end reasoning—seeking to choose effective means to one’s ends. Given this understanding of rationality, can there be limits to how rational it is to seek effective means to one’s ends? Can there be limits to how rational it is rational to be? I think so. There are times when means-end reasoning is counterproductive. If Tom is just getting involved in a new romantic relationship, for example, he might be well-advised to relax and let things happen. Thinking too hard about how to achieve his goals could make him look unromantic, and thus be precisely what stops him from achieving his goals.^{xvi}

Why be rational, then? This is not an idle question. It has an answer, namely, that being rational—involving as it does the seeking of effective means to our ends—generally serves our ends. This is not a tautology. It is an empirical claim. While we always have reason to endorse effective means to our ends, we do not always have reason to endorse means-end reasoning itself, for there are times when means-end reasoning cannot survive self-scrutiny, when means-end reasoning would be counterproductive. In such cases, means-end reasoning can lead us to reject means-end reasoning as a way of guiding action. Means-end reasoning is not guaranteed to pass its own test.

When means-end reasoning would not serve one's ends, then *choosing* to engage in means-end reasoning would not be rational by the lights of a means-end account of rational choice. Note that we can scrutinize means-end reasoning in terms of its means-end efficacy without questioning the underlying means-end account of rational choice. We can hold that rational choice involves seeking to choose effective means to one's ends and at the same time acknowledge, without inconsistency, that being in a rational-chooser frame of mind is not always to our advantage.

Further, to summarize earlier remarks, there are limits to the rationality of strategies and behavioral patterns commonly but contingently associated with rationality. A given strategy or persona is rational insofar as it serves the agent's ends, and sometimes the things that normally serve an agent's ends would be out of place. Rationality does not always dictate being strategic, or being calculating, or being self-controlled. It usually does not dictate being inhibited, and it almost never dictates being humorless.

VII. Making Room for Morality

As mentioned earlier, if we follow Gibbard and equate rational choice with what makes the most sense, all things considered, then our question about the difference between being rational and being moral resurfaces as a question about the difference between morality and prudence as domains within "all things considered" rational choice. For some purposes, it might also be useful to define a domain of epistemic rationality. We could, for example, distinguish epistemic from prudential rationality by saying the former concerns the pursuit of truth^{xvii} whereas the latter concerns the pursuit of one's long-range welfare.

While rational choice involves seeking to choose effective means to our ends, moral agency, as a first approximation, involves having certain ends and pursuing them within certain constraints. Conflict between rational choice and moral agency so conceived is not inevitable, but neither is convergence. We are left wanting to know which contingent

factors can get us from a means-end form of rationality to the substantive ends and self-imposed constraints of morality.

Prudential rationality counsels us to seek effective ways of serving our ends, and in particular our long range self-interest. Morality counsels us to do what is right. Each counsel seems incontestable on its own ground. From time to time, though, we need to choose. We can do what is in our best interest or what is right, but not both. It would be nice if we could settle which of the two—prudence or morality—has the stronger case, but the fact is that each has the stronger case on its own ground. Our only recourse is to explore the extent to which each, on its own ground, makes room for the other. Therein lies the strategy of this book.

Some people see rational choice as essentially tied not only to a formal notion of means-end efficacy but also to a substantive notion of self-regard. However, I see rational choice's tie to self-regard as merely contingent. It so happens that people tend to have predominantly self-regarding ends. Since our ends tend to be substantially self-regarding, rational choice tends to be aimed at self-regarding ends, but that is not the same as thinking those or any other substantive ends are an essential part of the concept of rational choice. There might be ways in which morality is coextensive with rationality without being coextensive in the same way with self-regard. Nor is morality exclusively tied to other-regard. Just as there are rational reasons for other-regard, so too (Chapter 8 argues) are there moral reasons for self-regard.

Identifying substantive grounds for moral endorsement is hard work reserved for later chapters, but without getting into questions about the content of moral reasons, let me say something here about their form. Deontology consists, in part, of a view that moral imperatives are categorical, which is to say their imperative force is not contingent on what the agent happens to desire. Against this, Philippa Foot says we “should be prepared to think that moral considerations give reasons for action only in ordinary ways” (1978, p. 154). When Foot speaks of giving reasons for action in ordinary ways, she means reasons

appealing to interests and desires. I am not at all sure that moral considerations give reasons for action only in ordinary ways. Still, only the ordinary ways are uncontroversial. Like Foot, then, I will explore the extent to which moral considerations do in fact give reasons for action in ordinary ways.

This approach does not presuppose that moral imperatives are like this: “If I want to be moral, I should keep promises.” Even if this imperative is true, it is still misleading at best as an account of why, morally speaking, I ought to keep promises. The reason why I morally ought to keep promises has little to do with whether I happen to want to be moral. (By way of comparison, consider: “If I want to be moral, then I *want* to keep promises.” We can agree that whether I want to keep promises can depend on whether I want to be moral, but this tells why I might want to keep promises rather than why I ought to.)

Hypothetical imperatives are imperatives such that, were we to put them in “if, then” grammatical form, their if-clauses would be couched in terms of appeals to the agent’s interests and desires. Hypothetical imperatives can appeal to a desire to be moral, but such an appeal makes them nonetheless hypothetical. Unlike Foot, I am perfectly content with the idea that moral imperatives are categorical. They present themselves to us as imperatives we ought to obey whether or not we want to. Morally speaking, we should keep promises, period, whether or not we want to be moral. The lesson to take from Foot’s argument is not that moral imperatives are hypothetical but rather that moral motivation is contingent. Even if moral imperatives are categorical, our motivation to obey them is hypothetical—contingent to some degree on what we desire.^{xviii} Later chapters, 5 and 6 especially, explore this contingency, asking how people can be led by ordinary interests and desires to internalize reasons for action that transcend ordinary interests and desires.

A related point: the categorical nature of moral imperatives is not what makes them moral. Etiquette and prudence also embody imperatives with a kind of normativity grounded in something other than an appeal to inclination. If we say X is a lie, or X is rude, or X is unhealthy, we are endorsing an abstention from X on grounds going beyond

appeal to occurrent desire. If we say that what Kate is doing is immoral (or is rude or is unhealthy) and Kate responds by saying “I’m grateful for your concern, but as it happens I don’t care about such things,” we will feel she has missed the point, for the normative force of what we said was not contingent on what Kate desires. On the contrary, the force is categorical.^{xix} Kate ought to care about such things, and not caring about them is shocking. Still, although any of these failures would be shocking in their own way, not all of them are moral failures. Although norms of etiquette are categorical, they have none of the special normative force of moral imperatives.

Morality’s special normative force—what makes it different from mere etiquette—is its teleology. Moral imperatives have a characteristic importance and urgency in virtue of being grounded in principles that are very general, nonarbitrary, and at the heart of what makes it possible for human beings to flourish in communities. Norms of etiquette are categorical, but to be categorical is not necessarily to be important. Norms of etiquette are in many ways quite like norms of morality, but they do not have the kind of central and ineliminably pervasive importance in human affairs that moral norms have.^{xx} The distinction, though, is largely one of degree. Really bad manners (like barging to the front of a lineup to buy tickets) shade into immorality. Really trivial immorality (like breaking a promise to bring dessert for a pot-luck party) shades into bad manners.

VIII. Conclusion

This chapter began by arguing that normative force—a reason for action—comes into being along with a person’s ends. When one has ends, one has reasons to act in some ways rather than others. Specifically, one has reasons to seek means to those ends. Along the way, we saw how this account makes room for a notion of objectivity. Indeed, without standards of objective rightness, there would be no point in trying to be subjectively

justified. A subjectively justified course of action is a means to an end, and whether it achieves its end is an objective matter.

My goal in characterizing rationality in means-end terms is not to be comprehensive but rather to start with a minimal account, uncontroversial as far as it goes, and see how far we can go with it. In the end, I will not deny that being rational involves more than seeking means to one's ends but neither will I start by assuming that it involves more than that.

One of the main reasons to be skeptical about means-end accounts of rational choice is that they seem to leave us with nothing to say about ends themselves. However, we will see that the means-end conception is not as obviously incomplete an account of rational choice as it might first appear to be. In Chapter 3, we will see that, for beings like us in situations like ours, rational choice extends to the evaluation of ends as well as means. This extension of rational choice is not a departure from the basic means-end concept, for it brings ends under the scope of rational choice by bringing in empirical facts about human nature rather than by expanding the concept *per se*.^{xxi} Surprisingly, the idea that means-end efficacy is a reason for endorsement will prove to be all we need, in terms of *normative* premises, to transcend instrumental rationality and develop a conception of rationality ranging over ends as well as means.

As one might infer from what has been said so far, my rational choice theory will be more closely tied to empirical considerations than is customary. My view is that decision theory has to be descriptive in part, because in order to have normative significance it has to engage both our purposes and our capacities as they really are. To have normative force for us, a model has to tell us about means to ends (selfish and otherwise) we actually have and it has to tell us about means that are actually available to us. What we ought to value depends on what we can value and on what would be good for us, which in turn depend on what we are like. The descriptive and normative facets of my theory will be intertwined throughout this book, because that is how it has to be for anyone aspiring to produce a normative theory relevant to beings like ourselves.

The first part of this book offers a theory about why we have the goals we do, and why we employ the strategies we employ—why we are not simple maximizers. More accurately, Part I offers a theory about why our goals and strategies are appropriate for beings like us. It is less concerned with the biological evolution of our goals and strategies and more concerned with why we have reason to be glad we have other-regarding ends and reason to be glad we are not relentless maximizers. We are better off as we are. Our self-regarding ends are better served as parts of a package containing other-regarding ends than they would be by a package that cut out the other-regarding elements.

With respect to strategies, one might presume doing as well as possible involves *aiming* to do as well as possible. I argued that this is not a tautology but a controversial empirical claim, sometimes true but often false. Chapter 2 argues that our goal of having life as a whole go as well as possible is best served by a combination of day-to-day strategies that often do not involve seeking optima.

ⁱ This refers to the question of whether we can validly deduce statements about what ought to be the case from statements about what is the case. See David Hume (1978, pp. 469-470) or G. E. Moore (1903, pp. 10-12).

ⁱⁱ As the phrase is used here, normative force is related to, but not identical to, motivational force. An end's normative force lies not so much in its tendency to move the agent as in the fact that being conscious of wanting some things rather than others gives the agent reasons to act in some ways rather than others. The distinction is adapted from J. David Velleman (1989, Chapter 7).

ⁱⁱⁱ I use 'end' and 'goal' (ditto for ought/should, decision/choice, and person/agent) interchangeably.

^{iv} What do we mean when we call a person rational? We are endorsing the person neither as a means to the person's ends nor as a means to someone else's ends. Typically we are saying that the person is using his or her cognitive capacities in a way that effectively serves his or her ends. Similarly, to call a choice rational is to endorse it as a means to the chooser's ends, but such endorsement implies an assumption that the choice involved deliberation. Otherwise, the endorsement is not apt. For more on conditions that support endorsement, and countervailing conditions, see the next section.

^v Robert Nozick advises: "Believe [hypothesis] *h* only if the expected utility of believing *h* is not less than the expected utility of having no belief about *h*" (1993, p. 92). Means-end efficacy is thus a necessary condition for rational belief in Nozick's theory, but not the only one. For example, hypothesis *h* must have at least minimal credibility and must be more credible than the alternatives.

^{vi} See Velleman (1989, Chapter 1) on the desire for self-understanding.

-
- ^{vii} Unfortunately, even if it is true that we strongly desire to justify ourselves to others on grounds they could not reasonably reject, it is also true that we desire simply to be accepted by others. Therein lies a source of psychological pressure that can undercut moral motivation, because there can be a big difference between what people do accept and what they should accept (or could not reasonably reject). Later chapters, especially Chapter 9, return to this point.
- ^{viii} I can only guess that I must have been influenced over the years by the work of my recently deceased and greatly missed friend and colleague John Pollock on defeasible reasoning. In conversation, John confirmed that what I call a supporting condition is what he would call a defeasible sufficient condition.
- ^{ix} Projects that begin by trying to analyze key concepts in terms of necessary and sufficient conditions tend to bog down in minutiae before they can get started. For most philosophical projects, supporting conditions are more to the point. I will rely on them again in Part II's analysis of moral agency.
- ^x Gibbard makes these remarks in a discussion of moral rightness, but he views rationality as encompassing all considerations of what it makes sense to do. I could be wrong, but the context strongly suggests that he intends his remarks to be about rationality in general.
- ^{xi} The same conclusion holds when a hiker knows that the river path will take him back to town but also knows that the river path is not the shortest path. In that case, the hiker maximizes expected utility not by seeking the path with maximum actual utility but rather by seeking a merely satisfactory path. See Chapter 2.
- ^{xii} There may also be a distinction between subjective and objective rationality. Insofar as a choice is rational if it warrants endorsement as a means to the chooser's ends, we can say a choice is subjectively rational if it is subjectively warranted and objectively rational if it is objectively warranted. I do not use this distinction in what follows, but it would hold up under any number of ways of drawing the prior distinction between subjective and objective. For example, we could distinguish between the best alternative as far as the agent knows and the best alternative in fact, although I think we express that distinction less ambiguously when we speak in terms of subjective justification versus objective success.
- ^{xiii} Chapters 2 and 4 explore this issue.
- ^{xiv} John Harsanyi (1982, p. 44) does equate rationality with maximizing expected utility, at least in the context of decisionmaking under conditions of risk or uncertainty, as does Richard Fumerton (1990, p. 101, 109). See also "Reason and Maximization" in Gauthier (1990).
- ^{xv} Also, as Michael Bratman (1990, p. 19) observes, we are "planning creatures. On the other hand, the world changes in ways we are not in a position to anticipate; so highly detailed plans about the future will often be of little use and not worth bothering with." We form partial plans and sometimes are appropriately casual about matters of detail.
- ^{xvi} This might be an example of a case where calling Tom's choice rational would not be a way of (even grudgingly) endorsing it. Calling Tom's choice rational might be our way of trying to explain to Tom

where he “went wrong.” We may be telling Tom that truly rational lovers realize that a rational-chooser frame of mind is not what is called for in a romantic context. I thank Jim Hamilton for this point.

^{xvii} It may be better to say the epistemic domain involves pursuing truths regarding specific questions that arise in the course of particular pursuits, rather than pursuing some more generic notion of Truth. On this view, the goals defining the epistemic domain have little content in abstraction from goals that define our pursuits in other domains and that raise particular issues with respect to which we want the truth. In this respect, then, goals defining the epistemic domain are parasitic on goals defining other domains.

^{xviii} Foot herself has come to regret her use of Kantian terminology (Foot, 1978, p. xiii). It turns out that she distinguishes between hypothetical imperatives and hypothetical ought-statements, which leads her to say rules of etiquette (and morality) may be hypothetical imperatives, yet nevertheless embody nonhypothetical uses of ‘should’ (1978, p. 161). I suspect, then, that Foot and I have no substantive dispute about the point expressed in the above text. We would agree about the sense in which moral imperatives are categorical.

^{xix} Or perhaps it would be better to say we are voicing an assertoric imperative. Assertoric imperatives (Kant, pp. 31-33) appeal to ends that human beings in general (perhaps even rational beings in general) are presumed to have. If a person denies having those ends, the denial would raise questions about the person’s rationality.

^{xx} Chapters 6 and 7 pursue this further.

^{xxi} Specifically, the extension takes account of the initial survival instinct and a certain capacity for reflection, as explained in Chapter 3.