

1 Morality and How We Ought to Live

Whereas primitive animals are ruled by their impulses, we humans are rational creatures who can respond to reasons for acting contrary to our impulses. We are therefore capable of directing our lives in accordance with our reasons rather than just being swept along by the flow of our impulses. So it's important for us to ask how we ought to live, as the answer to this question will tell us what living in accordance with our reasons entails. After all, how we ought to live just depends on how our various practical reasons – our reasons for living in different ways – combine and compete to determine *the way* (or, if there's a tie for first place, *the ways*) that we have most reason to live, all things considered. The relevant reasons are practical reasons. And *practical* reasons differ from other types of reasons in that they count for or against *performing (or intending to perform) certain actions*, whereas other types of reasons count for or against having some other kind of response. For instance, *epistemic* reasons typically count for or against *believing certain propositions*. And *evaluative* reasons typically count for or against *desiring certain states of affairs*. Now, I'll be concerned almost exclusively with practical reasons. So the reader should assume that, unless I explicitly state otherwise, the reasons that I'm discussing are practical reasons.

As just suggested, there seem to be different sorts of (practical) reasons. For instance, whereas the balance of moral reasons supports doing what's most morally choiceworthy, the balance of self-interested reasons supports doing what's most self-interestedly choiceworthy. But how do these different sorts of reasons combine and compete to determine which is most choiceworthy overall? And which is the way that we ought to live? Should we live as it would be most morally choiceworthy to live, as it would be most self-interestedly choiceworthy to live, or as it would be most overall choiceworthy to live? The question arises because these can seemingly come apart. After all, the most morally choiceworthy way to live may not be the most overall choiceworthy way to live. Consider that those known as *effective altruists* believe that the most morally choiceworthy way for you to live is by maximizing the good that you do over the course of your life.¹ Some of them even go so far as to suggest that you ought to earn as much as possible with the goal of giving most of what you earn (say, everything over the \$35,000 a year that you need to subsist) to effective charities. And they argue that it would be better for you to work on Wall Street and earn lots of extra money that you then give away to charity than to work in a profession in which you help people directly but make very little extra money that you can then donate to charity. They argue that, as a nurse, teacher, or social worker, you would make the world better only by the margin

¹ Two prominent effective altruists are William MacAskill (2015) and Peter Singer (2015).

that you do that job better than whoever would have otherwise done it. But as an effective altruist working on Wall Street, you can make much more of a difference because the person who would have otherwise filled your position would almost certainly have used the substantial extra earnings that you'll give to charity on luxuries for herself – and so you'll do much more good with those extra earnings than she would have.²

So, according to many effective altruists, the most morally choiceworthy life for a young person is the one in which they work to become an investment banker on Wall Street and then donate all but what they need for bare subsistence to the most effective charities that they can find. Such a person won't be able to afford to raise a family, for it isn't easy living in New York City on only \$35,000 a year. And they'll almost certainly need to live in a cramped apartment with several roommates just to afford the high rents in that city. What's more, they'll need to forgo expensive hobbies, such as golfing, scuba diving, music lessons, and art collecting. Last, they'll need to forgo the luxuries that many of us take for granted: vacations, dinners out, streaming services, nights at the theatre, etc. Indeed, when they get some time off work, they should probably use it to recover from a surgery in which they have donated one of their kidneys to a stranger in renal failure. For, as some effective altruists have argued, the risk of dying as a result of making such a donation is only one in four thousand and so to refrain from making such a donation would be to value their own life four thousand times more than that of the stranger.³

But even if living in the single-minded pursuit of doing the most good that one can is perhaps the most morally choiceworthy way to live, it may not be the most overall choiceworthy way to live. After all, it seems that we have good reason to pursue more well-rounded lives – see, for example, Berg (Manuscript). It seems, for instance, that I have good reason to live a life in which I raise a family, travel the world, attend the theatre, take music lessons, and pursue a career that I find more rewarding than any on Wall Street. Thus it seems that even if I could do more good by plugging away as an investment banker on Wall Street and donating most of what I earn to effective charities, I have good reason to want to interact face-to-face with the people whom I'm helping (e.g., my students) while doing the kind of work that I find fulfilling (e.g., teaching).

Of course all of this has been rather quick. I've just assumed that the effective altruist is correct about its being most morally choiceworthy to maximize the good that one does. And I've just assumed that the reasons that one has, say, to take a vacation or pursue a fulfilling career are nonmoral reasons, and that these

² See, for instance, MacAskill (2013). ³ For some examples, see Baggini (2015).

nonmoral reasons are not outweighed, overridden, or undermined by the moral reasons that one has instead to do as much good as possible. Yet it does seem that the most morally choiceworthy life is just the life that one has most moral reason to live. And it seems that the most overall choiceworthy life is just the life that one has most reason to live, all things (including both moral and nonmoral reasons) considered. So it seems that, in order to determine both what's most morally choiceworthy and what's most choiceworthy overall, we need to better understand what our reasons are, which of them are moral, which of them are nonmoral, and how they combine and compete to determine what we have most reason to do, all things considered. And this is what I aim to do, starting with an exploration of the nature of moral reasons.

2 The Nature of Moral Reasons

Moral reasons are a subset of normative reasons. Thus, to understand them, we must first understand what *normative* reasons are and how they differ from *nonnormative* reasons, such as explanatory reasons and motivating reasons. What's more, we need to understand how moral reasons differ from other kinds of normative reasons. In particular, we need to understand how *moral* reasons differ from *nonmoral* reasons. And we need to understand how normative reasons differ from facts that are relevant to how we ought to live but that are not themselves normative reasons – that is, how they differ from what I call *normatively relevant non-reasons*. The relations among these various notions are depicted in Figure 1.

2.1 Normative Reasons versus Nonnormative Reasons

A normative reason is a fact that counts for or against a subject's responding in a certain way to their circumstances. It counts for or against this response in that it counts for or against this being how they ought to respond.⁴ For instance, the fact that I'll probably need a root canal later if I don't get this cavity taken care of right away is a normative reason for me to get it taken care of right away. That is, it counts in favor of this being what I ought to do in that this will be what I ought to do unless some other fact defeats its favoring force. More precisely, then, we can define normative reasons for and against as follows.

⁴ Others believe that there's nothing to be said beyond saying that they count in favor of some response. Here's Tim Scanlon (1998, p. 17): "Any attempt to explain what it is to be a reason for something seems to me to lead back to the same idea: a consideration that counts in favor of it. 'Counts in favor how?' one might ask. 'By providing a reason for it' seems to be the only answer." And still others agree with me that more can be said but disagree with me about what more can be said. For instance, Julia Markovits (2014, pp. 2–3) holds that "what it *is* for a consideration to count in favor of an action is for it . . . to help satisfy one of [the agent's] desires."

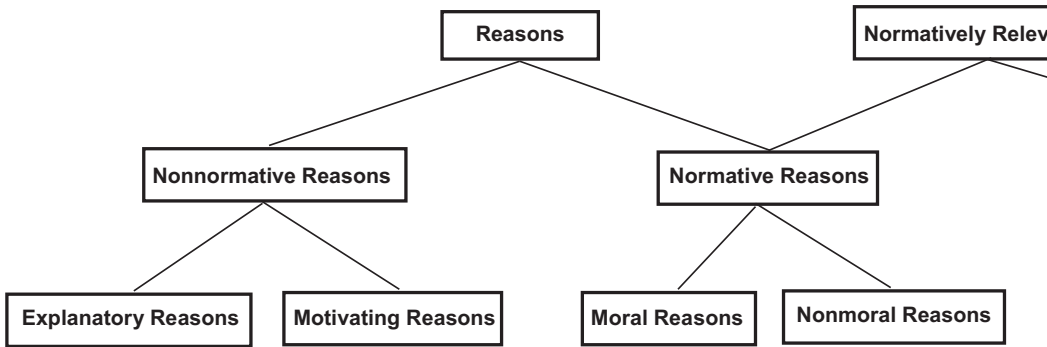


Figure 1 Reasons and Normatively Relevant Facts

Normative Reasons: For any subject *S* with the option of ϕ -ing in circumstances *C*, a normative reason *for* *S* to ϕ in *C* is any fact that will, absent its favoring force being defeated, make it the case that their ϕ -ing in *C* is what is most overall choiceworthy and, thus, what they *ought*, all things considered, to do in *C*.⁵ And a normative reason *against* *S*'s ϕ -ing in *C* is any fact that will, absent its disfavoring force being defeated, make it the case that they *ought not*, all things considered, to ϕ in *C*.

Given that there can be normative reasons both for and against one's ϕ -ing, normative reasons can conflict. The fact that I'll probably need a root canal if I don't get this cavity taken care of right away is a normative reason for me to get it taken care of right away, but the fact that I'm really pressed for time and need to finish a project at work by the end of the week is a normative reason against my doing so. Thus, whether I ought to take care of it right away depends (in part) on whether these reasons against doing so are defeated. And there are a number of possible ways that they might be. One possibility is that the reasons for getting the cavity taken care of right away *nullify* or otherwise *undermine* the disfavoring force of the reasons against doing so. To illustrate, suppose that I promised my wife that I wouldn't let work interfere with my healthcare. Perhaps this promise-related reason nullifies the work-related reason, undermining whatever disfavoring force that it would otherwise have. Another possibility is that health-related reasons always *override* work-related reasons such that even the weakest health-related reason for getting it taken care of right away trumps the strongest work-related reason against doing so. And a third possibility is that the reasons for taking care of it right away *outweigh* the reasons against doing so. For perhaps the combined favoring force of all the reasons for taking care of it right away outweighs the combined disfavoring force of all the reasons against doing so.⁶

⁵ Jonathan Dancy has claimed that some normative reasons for ϕ -ing (what he calls "enticing reasons" for ϕ -ing) favor one's ϕ -ing without being such that they will, absent being defeated, make it the case that one ought to ϕ . He claims that "they go to make a choice the best one, but not yet the one which one ought to take" (2004, p. 116). He denies that these enticing reasons can take us to an ought, because he denies that they can take us to an obligation and believes that we're obligated to do whatever we ought to do (2006a, p. 135). Yet it seems clear both that one ought to take whatever option one has most reason to take and that one isn't always obligated to take the option that one has most reason to take – more on this later in this Element. So, contrary to Dancy, I believe that what makes a normative reason for one to ϕ an enticing reason, rather than a requiring reason, is that it's a fact that favors one's ϕ -ing without giving anyone (even oneself) sufficient grounds for demanding that one ϕ s. It does, however, take us to an ought if undefeated.

⁶ When combining the forces of reasons, we should never combine the force of a derivative reason with that of the reason from which its force derives. That would involve a kind of problematic double-counting. Thus, when I talk about combining the forces of reasons, I'm talking about combining the forces of only nonderivative reasons. See Nair (2016).

Note, then, that the favoring and disfavoring forces of normative reasons can be more or less weighty. That is, normative reasons can favor or disfavor an action to a greater or lesser degree. And the weights of these forces can combine and compete with those of other normative reasons. What's more, the combined favoring force of a multitude of relatively weak reasons for one to ϕ could outweigh the disfavoring force of a relatively strong reason against one's ϕ -ing: the result being that one ought to ϕ even though none of the individual reasons that one has for ϕ -ing could, on its own, outweigh the strong reason one has against ϕ -ing. To illustrate, suppose that I have a weighty reason against going into the office today: I'll lose more than an hour of productivity just in the time it would take to commute to and from the office. And suppose that I have only one student who wants to meet with me today. In that case, the relatively weak favoring force of this reason would, we'll suppose, be defeated by the comparatively strong disfavoring force of the reason that I have against going into the office. But now suppose that I have instead a total of five students who each want to meet with me today. In that case, the combined force of each of these reasons would, we'll suppose, defeat the disfavoring force of the reason that I have against going into the office. So if I have only one student who wants to meet with me, I ought not to go into the office. But if I have five students who want to meet with me, I ought to go into the office.

As we've just seen, if I have an undefeated normative reason to ϕ in C, then ϕ is what I ought to do in C. And it's this ability to combine and compete so as to determine what one ought to do that's definitive of normative reasons. Normative reasons are normative in that they can contribute both to generating oughts and to preventing other reasons from generating oughts.

It's important to note, though, that a fact that doesn't provide an agent with a normative reason to ϕ in one set of circumstances could provide them with a normative reason to ϕ in another set of circumstances. For instance, the fact that I've promised that I would go into the office if and only if I have some other reason to go into the office doesn't provide me with a reason to go into the office in those circumstances in which I don't have any other reason to go into the office, but it does provide me with a reason to go into the office in those circumstances in which, say, I have a meeting at the office that I have good reason to attend (Dancy 2006b, p. 42).

It's also important to note that normative reasons, if not defeated, necessarily take us to an *ought*, but they don't necessarily take us to an *obligation*. For although we ought to do whatever we are obligated to do, not everything that we ought to do is something that we're obligated to do. To say that an act is one that you ought to perform is to say that it's the option that you have most reason to

perform, and it is therefore your most choiceworthy option.⁷ But to say that an act is one that you're obligated to perform is to say additionally that you would be blameworthy for failing to perform it absent some suitable excuse.⁸ Thus, an act can be one that you ought to perform without being one that you are obligated to perform. For you are not always blameworthy for failing to perform the option that you have most reason to perform – even when you lack a suitable excuse for so failing. Whether you are or not just depends on whether there is some actual or hypothetical person (either oneself or some other) who can rightfully demand that you perform this option. Thus, there has to be some (at least, possible) person who can rightfully hold you to account (by blaming you) if you fail to perform it without having some suitable excuse. To illustrate, consider the following claim: “I ought to send my mother flowers for Mother’s Day, but, at the very least, I must send her a card.” This claim, I take it, expresses the thought that what I have most reason to do is to send her flowers, but that I’m obligated at a minimum to send her a card. This is the least that I could do because she has every right to demand that I do at least this much given all that she’s done for me. And although neither she nor anyone else has the right to demand that I send her flowers, I ought to do so given what it would mean to her and how little it would inconvenience me to do so, which is what explains why this is what I have most reason to do.

As I’ve just explained, normative reasons are related to oughts; they can contribute both to generating an ought and to preventing other reasons from generating an ought. It’s potentially confusing, then, that in ordinary English we often use the word “reasons” to refer to things that are not related to oughts. So, to avoid confusion, we should label these reasons that are not related to oughts “nonnormative reasons.”

There are at least two types of nonnormative reasons: *explanatory reasons* and *motivating reasons*. I’ll consider each in turn. An explanatory reason explains why a subject responded as they did. For instance, the fact that I haven’t gotten enough sleep lately may explain why I’ve just snapped at you

⁷ I take all the following to be equivalent: “you ought to ϕ ,” “you have most reason to ϕ ,” and “ ϕ is your most choiceworthy option.” But I take all three to be distinct from “ ϕ is your best option.” For why, see Portmore (2019a, especially chap. 6).

⁸ Note that I’m assuming that there is nonmoral blame as well as moral blame – see Portmore (Forthcoming). What distinguishes moral blame from nonmoral blame is, I believe, the different sorts of reactive attitudes that are associated with each. For instance, moral blame (unlike nonmoral blame) involves in its first-person instances feelings of guilt. And this allows us to distinguish moral oughts and obligations from nonmoral oughts and obligations. If you fail to do what you’re morally obligated to do without any suitable excuse, it will be fitting for you to feel guilty. And if you fail to do what you morally ought to do without any suitable excuse, it will (or would be) fitting for you to feel guilty if there is (or had been) someone who could have rightly demand(ed) that you do as you morally ought to.

(Markovits 2014, p. 1). But of course my not having gotten enough sleep lately doesn't count in favor of my snapping at you. Rather, it's merely what caused me to snap at you. And so explanatory reasons are distinct from normative reasons – they are facts that explain why an agent did what they did without necessarily counting in favor of what they did. By contrast, a normative reason for performing an action, by conceptual necessity, counts in favor of performing that action.⁹

Another type of nonnormative reason is a motivating reason. Consider that even if I snap at you because I haven't been getting enough sleep lately, this won't be the consideration on the basis of which I chose to snap at you. Perhaps I chose to snap at you because you had been tapping your foot and this had gotten on my nerves (Markovits 2014, p. 1). But of course the consideration on the basis of which I chose to snap at you needn't be one that counts in favor of my doing so. The fact that you'd been tapping your foot doesn't count in favor of my snapping at you. Thus, motivating reasons are also distinct from normative reasons – they are the considerations on the basis of which the agent chooses to act without necessarily being considerations that count in favor of that act. And note that, unless I explicitly state otherwise, I'll use the term “reasons” to refer to normative reasons.

2.2 Qualified Normativity versus Unqualified Normativity

I've claimed that what's definitive of normative reasons is that they contribute both to generating oughts and to preventing other reasons from generating oughts. But what sorts of oughts? Interestingly, some philosophers deny that there's anything that we just plain ought to do. They claim instead that there is only what we legally ought to do, what we morally ought to do, what we self-interestedly ought to do, what we aesthetically ought to do, etc. That is, they deny that there is anything that we just plain ought to do and hold instead that there is only what we ought to do in various qualified senses. These philosophers are known as *normative pluralists*.¹⁰ To illustrate, suppose that I'm playing chess, I'm in check, and it's my turn.¹¹ The rules of chess dictate that, given that I'm in check, I must move my king one square to the left – this being the only move that gets me out of check. But the rules of etiquette dictate that I ought instead to resign by tipping over my king because I can no longer expect

⁹ Sometimes, when an agent is motivated to act by a fact that counts in favor of that act, that very fact counts not only as a motivating reason, but also as both an explanatory reason and a normative reason (Markovits 2014, p. 1).

¹⁰ Proponents of normative pluralism include David Copp (1997, 2009), Mathias Sagdahl (2014), and Evan Tiffany (2007). Critics include Owen McLeod (2001) and Dale Dorsey (2016).

¹¹ This example is borrowed from Richard Joyce (2001, p. 50).

a draw, let alone a win. Yet, from the point of view of self-interest, I ought to offer a draw because I don't want to lose and my opponent doesn't yet realize how untenable my situation has become. Last, let's suppose that morality directs me to violate the rules of chess by using my cell phone during the game given that I promised my wife that I would text her during the game. In such a case, normative pluralists hold that I morally ought to use my cell phone, self-interestedly ought to offer a draw, etiquetally ought to resign, and chess-rules-wise ought to move my king one square to the left. But they deny that there's anything that I just plain ought to do – that is, they deny that there's anything that I ought to do, all things considered.

We should, I believe, be skeptical of normative pluralism. After all, when ordinary people (those without advanced training in contemporary moral philosophy) use the word “ought,” they rarely, if ever, use any overt qualifier with it. That is, they rarely, if ever, use constructions such as “morally ought,” “legally ought,” “prudentially ought,” or “self-interestedly ought.”¹² And I doubt that they intend for there to be some covert qualifier attached to their “ought” that's supposedly fixed by the context. For it seems that when ordinary people use the word “ought,” they use it in an unqualified way that's meant to be authoritative such that if you don't respond as they say that you ought to, they take you to be guilty of a mistake from the point of view of reasons generally and not just from the point of view of some qualified normative perspective. Thus, we wouldn't tell someone that they ought to wait for everyone else to be served before taking a bite of their food when we know that they are a diabetic with a dangerously low blood sugar level. That is, we wouldn't presume that, in the context, our statement would mean only that this is what they etiquetally ought to do even though doing so would clearly be a mistake from the point of view of reasons generally.

Of course, whenever you violate a standard – whether it be one of legality, morality, etiquette, or reasons generally – you'll be guilty of the “mistake” of having violated that standard. But, for any standard other than that of reasons generally, it's an open question whether you'll also be guilty of the mistake of not responding appropriately to your reasons. After all, you may not have any reason to abide by the given standard. For standards come cheap in that they can

¹² On March 17, 2020, I did a search of the online version of the *New York Times* (www.nytimes.com). And whereas I got 820,437 results for “ought,” I got only two results for “legally ought,” one result for “morally ought,” and zero results both for “self-interestedly ought” and for “prudentially ought.” The one result for “morally ought” was due to a quotation from Oxford academic Richard Dawkins. And of the two results for “legally ought,” one seemed to be the result of a transcription error and the other was due to a quotation by a highly educated congressman who has both a BA and a JD from Harvard.

be constructed for any or no purpose at all.¹³ Consider, for instance, what we might call *the many-books standard*: an act meets this standard if and only if it involves touching three or more distinct books.¹⁴ Thus, the act of arranging a stack of several distinct books meets this standard, but the act of jogging around the neighborhood doesn't. Thus, in jogging around the neighborhood, one makes the "mistake" of violating the many-books standard. But if one doesn't have any reason to abide by this standard, then one isn't making the mistake of responding inappropriately to one's reasons.

So one reason to reject normative pluralism is the fact that it conflicts with the way that ordinary English speakers tend to use the word "ought." They tend to use it in an unqualified way such that those who fail to do as they say that they ought to are taken to be guilty of having failed to respond appropriately to their reasons.¹⁵ Indeed, few people would say that you ought to do something simply because it is in accordance with just any old standard, such as the many-books standard. For it is only some standards that we have reason to abide by. And because of this we can sensibly ask whether we just plain ought to do what we ought to do in some qualified sense – for example, whether we ought to do what we many-books-wise ought to do. So another reason to think that there is this unqualified sense of ought is the fact that we seem to be able to ask questions such as whether we ought to do what we legally, morally, or self-interestedly ought to do. And such questions presuppose that there is an unqualified sense of "ought."

Last, we should reject normative pluralism because it often seems to us that there's a normative question that needs answering even when we know what we ought to do in every qualified sense of "ought." For instance, sometimes I morally ought to do one thing but self-interestedly ought to do another. And yet it seems important to inquire further and ask what I ought to do, all things considered – that is, what I ought to do considering both what is most morally choiceworthy and what is most self-interestedly choiceworthy. And, in some

¹³ Some normative pluralists – those who we might call *restricted* normative pluralists (e.g., Copp 2009) – hold that there are only a limited number of domains that are normative in the relevant sense – perhaps just morality and self-interest. But they hold that, because these domains are incommensurable, there is nothing that we ought to do, all things considered, but only what we ought to do in the various qualified senses. A challenge for such restricted normative pluralists is to explain what the relevant sense of "normative" is such that some but not all standards of evaluation are normative if it's not just that some but not all such standards are those that we have unqualified reasons to abide by.

¹⁴ This is inspired by a similar example from Dale Dorsey (2016, p. 9).

¹⁵ As I see it, failing to do what one ought to do entails failing to respond appropriately to one's reasons in that it is inappropriate for one to fail to do what one has most reason to do. But if no one (not even oneself) can legitimately demand that one does what one has most reason to do in the circumstances, then one will be neither required to do so nor blameworthy for failing to do so.