

THERE IS NO MORAL OUGHT AND NO PRUDENTIAL OUGHT

Elizabeth Harman

1 Introduction

It is natural to think that there are a number of different *oughts*. There is a moral *ought*, there is a prudential *ought*, and so on. Furthermore, it is natural to think that each *ought* is such that one ought to do the best thing one could do, where the sense of *best* at issue varies with the kind of *ought* it is.¹ Thus, it is natural to think: morally, a person ought to do the morally best thing she could do, and prudentially, a person ought to do the thing that is best for her. One might also express these thoughts by saying: morality recommends that one do the morally best thing one could do, and prudence recommends that one do the thing that is best for oneself.

These natural thoughts suggest the further thought that the moral *ought* and the prudential *ought* often conflict, and thus that often, morally, one ought to do something although, prudentially, one ought to refrain from doing it.

While these thoughts are natural, and they express a commonly held view, I will argue that these thoughts are wrong. My modest aim is to show that there is an alternative view to the view I have described. My more ambitious aim is to show that my alternative view is correct. Once the two views are contrasted, I think it will be clear that although the commonly held view is indeed common, it is not supported by or warranted by ordinary moral thinking or ordinarily recognized moral phenomena, and we do better at capturing moral reality – and normative reality more broadly – by jettisoning the common view in favor of the alternative view I outline.

The alternative view denies all the natural thoughts I have described. It holds that there is no distinctively moral *ought*, though there are some *ought* facts that are distinctively moral. Similarly, there is no distinctively prudential *ought*, though there are some *ought* facts that are distinctively prudential. Finally, the alternative view holds that distinctively moral *ought* claims never conflict with distinctively prudential *ought* claims: it is never the case that, morally, one ought to do something, while prudentially, one ought to refrain.

2 The common view: there is a moral *ought* and a prudential *ought*

Consider the following view:

The Common View:

- (a) There is a moral *ought* such that: Morally, an agent ought to do something, just in case morality recommends that she do it. *and* Morally, an agent ought to do something, just in case it is the morally best thing she could do. *and* Morality recommends that an agent do something just in case it is the morally best thing she could do.
- (b) There is a prudential *ought* such that: Prudentially, an agent ought to do something, just in case prudence recommends that she do it. *and* Prudentially, an agent ought to do something, just in case doing it would be best for the agent. *and* Prudence recommends that an agent do something just in case doing it would be best for the agent.
- (c) These two *oughts* often conflict: it is often the case that morally, one ought to do something, while prudentially, one ought to refrain from doing it.

I will argue that the Common View can be (and should be) rejected. I will present an alternative picture. On the alternative picture, it is indeed sometimes true that morally, one ought to do something, and it is sometimes true that prudentially, one ought to do something, but these truths do not involve a distinct moral *ought* nor a distinct prudential *ought*. The alternative picture embraces the first biconditional under (a), embracing the tight connection between what, morally, one ought to do and what morality recommends. And the alternative picture embraces the first biconditional under (b), embracing the tight connection between what, prudentially, one ought to do and what prudence recommends. But the alternative picture rejects the other two conditionals under (a) and the other two conditionals under (b): on the alternative picture, the asserted tight connections with the morally best thing the agent could do, and with what would be best for the agent, do not hold.

The plan for the rest of chapter is as follows. In sections 3, 4, and 5, I discuss and argue against the Common View, as it concerns morality. Sections 6 and 7 address objections. In section 8, I present the part of the alternative view that concerns morality. In section 9, I argue against the Common View, as it concerns prudence. Section 10 presents the remaining part of the alternative view, concerning prudence. Section 11 discusses some objections. Section 12 summarizes the chapter's conclusions. And, finally, section 13 explains the broader significance of these conclusions.

3 Moral *ought* statements do not just state moral requirements

In this section, I will explain why people may have been drawn to the Common View. To see the appeal of the Common View, let's contrast it with another view. The Common View includes the following claim:

- (★) Morally, an agent ought to do something, just in case it is the morally best thing she could do.

The contrasting view is:

The Naïve View: Morally, an agent ought to do something, just in case she is morally required to do it.

I call this view "naïve" because I don't think any non-consequentialist holds it. (It would, however, be endorsed by a maximizing consequentialist.)

These two views agree about cases involving morally required options. For example, consider this case:

Aaron promised Bill that he would go to Bill's poetry reading. Aaron does not enjoy poetry, but Bill is his friend, and it will mean a lot to Bill to have Aaron there. On the day of the reading, Aaron is invited to join another friend at a Lakers game. Aaron loves the Lakers, though of course there are other games he could attend this season. Aaron tells all this to Carl. Carl says, "Sorry! I know it's tempting to bail on Bill, but morally, you ought to go to the reading."

In this case, it is true that, morally, Aaron ought to go to the reading.² Both the Common View and the Naïve View accommodate this truth, because going to the reading is Aaron's best option, and it is also a morally required option. Similarly, consider this case:

Donna is an amateur tennis player who maintains a fierce and genuinely hostile rivalry with Ellen, another tennis player. Donna is doing a preliminary review of applications for a job at her company, after which a committee will make a decision among the strongest candidates. Donna sees that Ellen's application is very strong. Donna hates the idea of working with Ellen and knows that she could eliminate Ellen's application from consideration. Donna tells all this to Fiona. Fiona says, "Sorry! I know it's tempting to reject Ellen at this point, but if her application is strong, then morally, you ought to keep her under consideration along with the other strong applicants."

In this case, it is true that, morally, Donna ought to keep Ellen under consideration.³ Both the Common View and the Naïve View accommodate this truth, because keeping Ellen under consideration is both Donna's best option and it is morally required. The same will hold for any case involving a morally required option. Whenever an agent has a morally required option, then that option is also her morally best option, and so both the Common View and the Naïve View will imply, correctly, that morally, she ought to take that option.

To see why one would reject the Naïve View in favor of the Common View, let's turn to considering some cases involving supererogatory actions. In ordinary life, we often say to each other, truly, "You *ought to do it*, but you don't *have to do it*," where by "you don't have to do it," we mean that it isn't morally required, and in saying "you ought to do it," we are making a *moral claim*. For example, consider the following case:

Georgia's elderly neighbor Harriet is recovering from surgery. Georgia knows that, given her friendship with Harriet, she is morally required to visit her at least a few times a week, and Georgia has already visited Harriet several times this week. Today Georgia comes home from work and is a bit tired but sees Harriet's light on across the street. Georgia could stop by to see Harriet for ten minutes. It would lift Harriet's spirits and wouldn't cost Georgia very much; she'd just start cooking her dinner ten minutes later. The morally best thing Georgia could do at this moment is to go spend ten minutes with Harriet.

In this case, it is natural to think that Georgia is not morally required to spend ten minutes with Harriet, but that Georgia *ought to spend ten minutes with Harriet*. It is natural to think that, morally, Georgia ought to spend ten minutes with Harriet. It seems that *morality recommends*

that Georgia spend ten minutes with Harriet. This thus seems to be a case in which morality offers a *recommendation* that is not a *requirement*. The Naïve View does not allow that moral *ought* statements can do this. According to the Naïve View, morality recommends all and only those actions that are morally required. By contrast, the Common View can accommodate all of these claims: spending ten minutes with Harriet is Georgia's morally best option at the moment, and so, morally, she ought to take it, according to the Common View.

Consider another case of a supererogatory action, which also illustrates this point:

It is November 2016. Isaac is a college student who is worried after Donald Trump's election to the U.S. presidency. He donates his limited spare money to Planned Parenthood and the ACLU. He hears that some Muslim students are afraid of hate crimes and that some students are volunteering to walk them home from prayer at the local Mosque. It's raining and Isaac is tired. But this is a concrete way he could help. "I ought to go volunteer. I don't have to, but I ought to do it," he thinks.

Isaac is making a moral claim, and it is a moral claim that might well be true. The Naïve View cannot accommodate the truth of Isaac's *ought* claim, because volunteering is not morally required. But the Common View can accommodate the truth of his claim: morally, Isaac ought to volunteer.

The cases of Georgia and Isaac show that the Naïve View is false. Sometimes, morally, one ought to do something, although it is not morally required. These cases support the Common View instead.⁴

The Common View is indeed common; I hypothesize that philosophers have embraced the Common View by overgeneralizing from cases like Georgia and Isaac.

Let's take stock. The Common View includes these two claims:

- (★) Morally, an agent ought to do something, just in case it is the morally best thing she could do.
- (★★) Morality recommends that an agent do something just in case it is the morally best thing she could do.

As we've seen so far, true instances of these claims are given by morally required actions, such as Aaron's going to the reading and Donna's keeping Ellen under consideration, and by some supererogatory actions, such as Georgia's spending ten minutes with Harriet and Isaac's volunteering to walk the Muslim students home. In the next section, I will argue that some instances of claims (★) and (★★) are false, and thus that the Common View must be rejected.

4 Sometimes (often), morality does not recommend that one do the morally best thing one could do

Consider this case involving a supererogatory option:

James is an accountant who paints paintings as a hobby. He has always wanted to have a show in a professional gallery but has not ever had one. The local gallery offers a competition: artists can submit one piece, and if their piece is selected, they get a show for one week in the gallery. The gallery is a popular destination for other gallery-owners and art buyers, so a one-week show there might lead to further successes as

a professional artist. Getting the show would be the accomplishment of a long-held dream for James and might lead to further meaningful accomplishments. James works very hard on his piece to submit. On the day of the deadline, James is getting in his car to drive to the gallery, with the deadline in half an hour. James's neighbor Kenny, who is seven years old, wanders over and asks if James will play *Parcheesi* with him. James and Kenny do play *Parcheesi* sometimes. Kenny's mom has been sick, and Kenny is often sad and lonely, so it means something to Kenny when James plays with him. James is about to leave town for a work trip, so this would be his last chance to play with Kenny for a while. James has two options.

- (i) James could play *Parcheesi* with Kenny, missing his chance to submit to the art show.
- (ii) James could tell Kenny he can't do it today and go on to the art gallery.

In this case, what is James's morally best option? Well, staying and playing with Kenny would be a morally good thing to do. But going to the art gallery, while there is a lot to be said in its favor, is not a *morally* good thing to do. And whenever one option is morally good but an alternative is not morally good, then the first option is morally better. Because these are James's only two options, it follows that playing with Kenny is James's morally best option.⁵ But it would be a big mistake for James to stay to play with Kenny. This is a case in which taking his morally best option would involve James making a big mistake.

Now, suppose that morality always recommends that one take one's morally best option, as claim (***) says. If so, then this is what morality would say to James: "James, play with the neighbor boy." or "James, we recommend that you stay to play with Kenny, though you don't have to do that." Surely morality does not say this. Only a jerk would give James this advice. The advice that morality gives is not advice that only a jerk would give. Thus, this case shows that claim (***) is false. Furthermore, what morality recommends is exactly what, morally, an agent ought to do. Thus, the case also shows that claim (*) is false. It is not true that, morally, James ought to stay to play with Kenny. Thus, it is not always true that, morally, one ought to take one's morally best option; claim (*) is false.

Consider the following case, which also involves a supererogatory option:

Laura is a distinguished professor who has had a long day at work, at the end of which she has a meeting with Mark, a young man in his mid-twenties who is seeking her advice on pursuing a career in her field. Toward the end of the meeting, Laura mentions that she has a six-month-old baby. Mark says "Oh, so you're not really writing right now!" Laura can tell that he is attempting to be friendly and to have a moment of human connection, but she's offended and annoyed by his remark. His remark is sexist and, in her case, inaccurate. Laura has several options at this moment.

- (i) She could kindly take the time to explain to Mark why what he said is problematic, thus perhaps saving him from making future comments that might hurt him professionally and thus potentially saving some women from being the recipients of such comments, while offering the explanation in such a nice way that she does not make Mark feel too bad.
- (ii) She could proceed as though she is not at all offended and end the meeting on a friendly note.

- (iii) She could reveal her annoyance, ending the meeting politely but not warmly, and get back to her work.
- (iv) She could express her annoyance, pointedly and sharply, not sparing Mark's feelings.

Those are Laura's four options. Laura is tired after a long day and does not feel like spending her limited mental energy on educating yet another naïve young man. She takes option (iii), ending the meeting politely but not warmly.

In this case, all four of Laura's options are morally permissible. But option (i) is her morally best option: it would be a morally good and kind thing for her to do, though it is not morally required. Nevertheless, morality does not *recommend* that Laura take option (i). If morality did make this recommendation, morality would be saying to Laura, "Laura, educate the young man in your office." or "Laura, although you are tired, and you have to deal with sexist stuff like this all the time, we recommend that you take the time to kindly explain to this man why what he said is problematic, and do it in a way that spares his feelings as much as possible." Only a jerk would say this to Laura. The recommendations of morality are not recommendations that only a jerk would make. So this case shows that claim (***) is false: morality does not always recommend that one do the morally best thing one could do. Furthermore, what morality recommends is exactly what, morally, an agent ought to do. Thus, it is not the case that, morally, Laura ought to take option (i). So this shows that claim (*) is false: it is not always true that, morally, an agent ought to take her morally best option.

In this section, I've argued for the strong conclusion that the Common View is false. But we can also take this section as arguing for the weaker conclusion that a rejection of the Common View has some plausibility. If I haven't convinced you of my claims about the cases of James and Laura, I hope to have begun to show what an alternative view would look like and why it has some plausibility. I will elaborate on that alternative view in section 8.

5 Sometimes morality's recommendation is to do a morally good thing that is not the agent's morally best option

In section 4, I argued that certain cases of the following kind provide counterexamples to the Common View: cases in which an agent has a morally good option, but morality does not recommend that the agent do something morally good. James could play with Kenny rather than submitting to the art contest, and Laura could kindly explain the sexist remark rather than end her meeting abruptly; these are morally good options available to these agents. Yet in these cases, there is no morally good option that is recommended by morality. In these cases, there is no morally good option such that, morally, the agent ought to take that option.

In this section, I will argue that another kind of case also provides a counterexample to the Common View: cases in which morality does recommend that the agent do something morally good but does not recommend her morally best option. In these cases, there is a morally good option such that, morally, the agent ought to take that option. Nevertheless, these cases provide counterexamples to the Common View.

Consider the following case:

Nicole is a waitress at a restaurant where some of her regular customers are Deaf and communicate with each other in sign language. She wishes she knew sign language

so that she could communicate with them in their primary language, in order to be friendly and welcoming on a regular basis. She looks into a class in American Sign Language, but the class would be expensive and the time it would take would cause her to lose valuable sleep. Instead, Nicole could study some YouTube videos that teach basic ASL signs. It would take some work, but she could learn enough to be able to greet her Deaf customers in ASL. Nicole thinks to herself, “I ought to study these YouTube videos, though I don’t have to.”

Nicole’s claim is a moral claim, and it may well be true. Let’s focus on a version of this case in which Nicole’s claim is true. Morally, Nicole ought to watch the YouTube videos. This is a morally good thing that she could do and, morally, she ought to do it. But watching the YouTube videos is not Nicole’s morally best option. The morally best thing she could do in this case is to take the ASL class. Taking the ASL class would be an even nicer thing for Nicole to do for her Deaf customers. But even though this is the morally best thing she could do in this case, it is simply not true that, morally, she ought to take the ASL class. Furthermore, it is not true that morality recommends that Nicole take the ASL class. Instead, morality recommends that Nicole study the YouTube videos.

This is a case in which Nicole has a morally good option available – watching the YouTube videos – which is recommended by morality and which is such that, morally, she ought to take it. Nevertheless, this case shows that claims (*) and (**) are false, because Nicole’s morally best option – taking the ASL class – is not such that, morally, she ought to take it and is not such that morality recommends that she take it.⁶ Thus, this case shows that the Common View is false.⁷

If you’re not convinced by what I say about Nicole’s case, we can take her case as establishing the weaker claim that there is some plausibility to a denial of the Common View;⁸ I elaborate on an alternative view in section 8.

6 Objection: the moral perspective ignores self-regarding reasons

Someone might object to the claims I’ve made about the cases of James, Laura, and Nicole as follows:

Whether, morally, one ought to do something is of course different from whether one ought to do it, all things considered. It is not true that Laura or James, in their situations, ought to engage in helping behaviors, all things considered. But morally, that is what they ought to do, because that is the morally best thing that each can do. When we ask what an agent ought to do, morally, or (equivalently) what is the morally best thing she could do, or (equivalently) what morality recommends that she do, we must remember that morality takes a certain perspective. Consider the fact that James would be making a big sacrifice in staying to play with Kenny. This fact counts against staying, when we consider what James ought to do, all things considered. But it simply doesn’t count that way when we are asking what the morally best thing James could do is; thus, it simply doesn’t count that way when we are asking what James ought to do, morally.

This objector holds that the Common View is correct and that my claims simply misunderstand the nature of the moral perspective, from which morality’s recommendations are issued. The objector makes the following claim:

An agent’s self-regarding (or self-interested) reasons against acting in certain ways simply don’t count against acting in those ways from the moral perspective. That is, that

an action would be burdensome to an agent simply doesn't count against doing it, from the perspective of morality.

But this claim is obviously false. Morality does see the force of self-regarding reasons against action.⁹ Morality takes self-regarding reasons into account in failing to require certain actions, and in failing to prohibit certain actions, because of burdens that agents would otherwise be morally required to endure. For example, one is typically morally required to keep one's promises, but if keeping a promise would prove unexpectedly burdensome, breaking the promise is often morally permissible. If morality recommended to Laura and James that they perform the helping actions available to them, or if morality recommended to Nicole that she take the ASL class, then morality would be ignoring some considerations to which morality is otherwise sensitive.

7 Objection: something can be good advice although only a jerk would offer it

In section 4, I argued that James and Laura both have supererogatory options that are the morally best things they could do but that it is false that, morally, they ought to take these options. I appealed to the fact that only a jerk would say to either James or Laura, "I recommend that you do this, despite the cost to you of doing it." An objector might respond as follows:

Sometimes it is not okay to offer certain advice to someone, even though it's good advice. Sometimes only a jerk would urge someone to do something that she is morally required to do. Sometimes only a jerk would urge someone to do something that, morally, she ought to do.

I will offer two responses to this objection.

First, I want to clarify my attitude to the argument in section 4. In that section, I do indeed argue *from* the claim that only a jerk would say "I recommend that you do this" to the claim that it's false that morally, the agent ought to take that option. I do take that consideration to support that conclusion. But I also think that it is independently plausible that it is false that, morally, James ought to stay to play with his neighbor, and it is independently plausible that it is false that, morally, Laura ought to kindly explain to the young man why his comment is sexist, sparing his feelings as much as possible.¹⁰ So, while I do argue for these moral claims in section 4 – and then go on to use these moral claims to argue that the Common View is false – ultimately I don't want all the argumentative weight of that section to fall on my claim about what follows from the fact that only a jerk would recommend these actions; I think my moral claims about these cases are independently plausible, and I'm happy to rely on them directly, in arguing against the Common View.

Having said that, let me now turn to defending the move *from* the fact that only a jerk would say "I recommend that you do this" in these cases *to* the claim that it is false that, morally, each agent ought to take their supererogatory option. My second response to the objector is as follows. I agree with the objector that sometimes only a jerk would recommend to someone that she do the thing that she is morally required to do. And I agree with the objector that sometimes only a jerk would recommend to someone that she do the thing that, morally, she ought to do. But in what kinds of cases is this true? Suppose that someone is morally obligated to do something difficult; she is trying to do it, but she's not really succeeding; she is clearly conscious of her obligation; and no good would come from urging her to do what she is already trying to

do, except to make her feel terrible. For example, if one saw one's friend struggling to be patient with her difficult five-year-old son, no good would come of recommending that she be patient; only a jerk would do this. Or suppose that someone faces a difficult choice over whether to do the right thing; he has struggled over his choice, and decided with some regret to violate his obligation; suppose a friend knows that she cannot persuade him to change his mind and that recommending the right action would do no good. For example, if one's friend has decided not to confess to a minor fraud at work, the confession of which would cause the financial ruin of his family, leaving them destitute, then although he has chosen wrongly, it may be that only a jerk would say "turn yourself in." These are indeed two kinds of cases in which only a jerk would recommend to someone that she do the thing that, morally, she ought to do. These are cases in which the agent already knows that, morally, she ought to do it, and she is either struggling to do it, or she has already decided for sure against doing it. These cases are very different from the cases of James and Laura. Consider James in particular. He is not struggling with a difficult choice. And, if he starts saying "no" to playing with Kenny, James is not making any kind of mistake. My response to the objector is that while it is true that *sometimes*, morally, a person ought to do something, although only a jerk would recommend that action, this phenomenon arises in a particular kind of case, and the cases of James and Laura are not this kind of case. When we see that only a jerk would recommend that James or Laura take their morally best option, we have no reason to think that *nevertheless* that option is recommended by morality.¹¹

8 Making sense of *morally, you ought to do it*

In this section, I will outline the first part of my proposed alternative view, on which the Common View is false. (The first part of My Proposed View concerns morality; the second part of My Proposed View concerns prudence.) The view I develop acknowledges that it is sometimes true that, morally, a person ought to do a particular thing and acknowledges that some *ought* facts are *moral facts*, as in the cases of Aaron, Donna, Georgia, Isaac, and Nicole. Here is my alternative proposal:

My Proposed View – Part 1:

An agent is such that, morally, she ought to ϕ *just in case* all things considered (in light of all of her reasons), she ought to ϕ *and* there are some moral considerations in favor of ϕ ing that centrally explain its being the case that she ought to ϕ .

It follows from My Proposed View that every morally required option is one that, morally, the agent ought to take. This follows because if an option is morally required, then the agent ought to take it, all things considered. (In saying this, I am assuming that moral requirement is overriding.) And if an option is morally required, then there are some moral considerations in favor of taking it that centrally explain its being the case that the agent ought to take it, all things considered. So, on My Proposed View, it is true that, morally, Aaron ought to go to the poetry reading, and it is true that, morally, Donna ought to keep Ellen under consideration.¹²

The cases involving supererogatory options that I've discussed in this chapter show that the following claim is true:

The Supererogation-Ought Claim:

Some (but not all) supererogatory options are such that, all things considered, the agent ought to take them. (And some of these are such that the moral considerations that make them morally good to take also centrally explain its being the case that they ought to be taken, all things considered.)

Among moral philosophers who acknowledge the existence of the supererogatory, there has not been adequate appreciation of the fact that sometimes an agent has one or more supererogatory options and the agent ought to take one of them, while sometimes an agent has one or more supererogatory options and yet it is not true that the agent ought to take one of them. The cases of Georgia, Isaac, and Nicole exhibit the first phenomenon: Georgia ought to visit her neighbor for ten minutes, Isaac ought to volunteer to walk the Muslim students home from prayer, and Nicole ought to study the YouTube videos; these are supererogatory options that the agents ought to take. (My considered view about these cases is not that the described details of the cases settle that these *ought* claims are true. Rather, I take the more modest view that the details of these cases are *consistent with* the truth of these *ought* claims. My view is simply that there are versions of the cases of Georgia, Isaac, and Nicole in which these *ought* claims are true.) The cases of James and Laura exhibit the second phenomenon: while both James and Laura have supererogatory options available to them, neither of them ought to take a supererogatory option. (My considered view is simply that there is a version of Laura's case in which it is not the case that she ought to take her supererogatory option.)

The Supererogation-Ought Claim is true because moral considerations continue to have force within the realm of the morally permissible; they can settle that, all things considered, one ought to do something, although it is not morally required. This fact has not been appreciated by philosophers.¹³ Taking My Proposed View and the Supererogation-Ought Claim together, the following claim follows:

The Supererogation-Morally-Ought Claim

Some supererogatory actions are such that, morally, the agent ought to perform them.
Some of these are the agent's morally best option; some are not.

The cases of Georgia and Isaac are cases in which, morally, an agent ought to take a supererogatory option which is their morally best option. The case of Nicole is a case in which, morally, an agent ought to take a supererogatory option, but it is not her morally best option.

On My Proposed View, while sometimes it is true that, morally, an agent ought to take her morally best option, in many cases, an agent's morally best option is not such that, morally, she ought to take it. We saw this in the cases of James, Laura, and Nicole. None of these agents is such that, morally, they ought to take their morally best option. Furthermore, in many cases, there is no option, not even a disjunctive option, such that, morally, the agent ought to take that option. We saw this in the cases of James and Laura. The options that James and Laura ought to take, all things considered, are not centrally supported by moral reasons.

9 Prudence does not always recommend the option that is best for the agent

In this section, I argue against the part of the Common View that is about prudence. My argument is parallel to my argument about morality, and the alternative view I develop is parallel.

The Common View makes the following claims:

- (★★★) Prudentially, an agent ought to do something, just in case doing it would be best for the agent.
- (★★★★) Prudence recommends that an agent do something just in case doing it would be best for the agent.

There are plenty of cases in which prudence does recommend taking the option that would be best for oneself. Suppose a pipe has burst in your basement, and it is spewing water everywhere and starting to flood the basement. The best option for you is to go down there right away, get soaked in water, and put a stop to the flooding as quickly as possible; it will be worse for you if you delay. And indeed, prudence recommends that you go down right away. Suppose you have a regular workout schedule in which you go to the gym every Monday morning. You wake up on Monday morning, and are deciding whether to go to the gym. The option that would be best for you is to go to the gym, and indeed prudence recommends that you go to the gym. In each of these cases, the option that is best for the agent is indeed such that, prudentially, she ought to take it.

However, there are cases in which it is not true that prudence recommends performing the action that is best for the agent. Consider Aaron, who is deciding between going to the poetry reading he promised to attend and seeing a Lakers game he would greatly enjoy. The option that is best for Aaron is going to the game. However, prudence does not recommend that Aaron go to the game. Prudence doesn't recommend that Aaron fail to act as morality requires. It's just not true that, prudentially, Aaron ought to go to the game. Consider Isaac, who is deciding whether to volunteer to help keep Muslim students safe, which would involve going out in the cold rain. The option that is best for Isaac is staying warm and dry inside, but prudence does not recommend that Isaac stay inside, and it is not true that, prudentially, Isaac ought to stay inside. Prudence doesn't recommend that Isaac refuse to do something to help others, even at some cost to himself.

Consider also the following case, which is a fictionalized version of real events (involving a non-philosopher perpetrator):

Three students have accused a famous professor, Professor X, of sexual harassment. Professor X has responded by suing the students for defamation. The case is widely reported in the news. Another professor says the following about the case: "It is exceedingly unlikely that these three students are lying about what happened to them. So Professor X's lawsuit against them seems to be an effort to intimidate them, though a risky one, as the lawsuit may well dig up other aspects of his past. Given that Professor X's academic career is over unless he prevails, taking this gamble is something that, prudentially, Professor X ought to do, even if he's guilty."

What should we make of these remarks? The speaker claims that prudentially, Professor X ought to sue the students and thus that prudence recommends that Professor X sue the students. But surely both of these claims are false. Notice that the speaker takes it to be exceedingly likely that Professor X has sexually harassed these three students and is now subjecting them to a difficult and stressful defamation lawsuit *on top of* having harassed them. This is the kind of thing that has a high likelihood of driving these students out of their chosen field of study, and some likelihood of driving them into depression or suicide. Surely it is not true that, prudentially, one ought to continue to abuse the people one has sexually harassed by suing them for getting up the courage to tell the truth about what one did. Surely prudence does not recommend such a thing. Prudence does recommend self-preserving activities like eating right and going to the gym. Prudence does not recommend deeply immoral behavior like suing the women that one has sexually harassed to try to scare them into withdrawing their complaints.

The cases of Aaron, Isaac, and Professor X are cases in which prudence does not make any recommendation at all, although the agent has an option that is better for the agent than an alternative. There are also cases in which prudence *does* recommend performing an action that

is good for the agent, and yet prudence does not recommend performing the action that would be best for the agent. Consider the following:

Olive will fail out of school unless she passes tomorrow's test, but she will have to miss the best party of the year if she spends tonight studying. She knows how she could cheat on the test in a way that would not be discovered. Olive could (i) go to the party and cheat on tomorrow's test, (ii) spend tonight studying and miss the party, or (iii) go to the party and then fail the test.

In this case, the option that is best for Olive is option (i), go to the party and cheat on the test. This would enable her to both pass the test and enjoy the party. However, prudence does not recommend that Olive take option (i). Rather, prudence recommends that Olive study tonight. Prudentially, Olive ought to study tonight. This is a case in which prudence makes a recommendation, but it does not recommend the option that is best for the agent.

The cases of Aaron, Isaac, Professor X, and Olive show that the Common View is false. Or, more modestly, my discussion of these cases shows that an alternative view of prudence can be held: if you do not agree with my claims about these cases, I hope they start to show what an alternative view to the Common View would look like and why it has some plausibility. In the next section, I spell out the alternative view.

10 Making sense of *prudentially, you ought to do it*

We can deny the Common View while acknowledging that it is sometimes true that, prudentially, an agent ought to take a particular option.

My Proposed View – Part 2:

An agent is such that, prudentially, she ought to ϕ *just in case* all things considered (in light of all of her reasons), she ought to ϕ *and* there are some prudential considerations in favor of ϕ ing that centrally explain its being the case that she ought to ϕ .

According to My Proposed View, sometimes an option that is best for the agent is such that, prudentially, the agent ought to take that option. (We see this in the case of going to the gym.) But often the option that is best for the agent is not recommended by prudence, and it is not true in such cases that, prudentially, the agent ought to take that option. (We see this in the cases of Aaron, Isaac, Professor X, and Olive.) In many cases, there is no option, not even a disjunctive option, such that, prudentially, the agent ought to take that option. (We see this in the cases of Aaron, Isaac, and Professor X.)

11 Objections from other *oughts*

In this section, I will consider two objections from consideration of the existence of other purported *oughts*. The first objector addresses me as follows:

Consider the following truths:

Grammatically, one ought to refrain from ending a sentence with a preposition.
Football-wise, the cornerback ought to tackle the receiver when he is running with the ball.

Etiquette-wise, one ought to use a smaller fork for salad and a larger fork for one's main course.

Legally, one ought to carry one's driver's license whenever one drives a car.

It looks as though there are quite a few *oughts*: an *ought* of grammar, one of football, one of etiquette, one of the law, and so on. Here is a dilemma: either you hold that there are not any distinctive *oughts* like this and apply your account generally to all apparent distinct *oughts*, or you hold that the purported moral *ought* and the purported prudential *ought* are special. If you take the first horn of the dilemma, you face counterexamples: football, etiquette, and the law sometimes recommend actions that one should not take, all things considered. (For example, given the risk of concussions, both to oneself and the other player, one should probably never tackle anyone.) If you take the second horn of the dilemma, you face the burden of explaining why the apparent moral and prudential *oughts* differ from these other *oughts*.

I choose the second horn of this dilemma. Morality and prudence are special. Unlike football, grammar, etiquette, and the law, morality and prudence have a particular feature:

Morality does not give bad advice.

Prudence does not give bad advice.

This distinguishes morality and prudence from the other sources of advice.¹⁴

Would I generalize my arguments, and My Proposed View, to these other purported *oughts*? I would not, simply because it is not at all plausible that these other sources of advice never give bad advice.

Consider now another objection:

It may sound strange to say that *prudentially*, Aaron ought to ditch Bill's poetry reading and go to the Laker game. But it sounds fine to say that *self-interestedly* or *selfishly*, that's what Aaron ought to do. But since prudence is simply the same as self-interest, this shows that it's actually true that, prudentially, Aaron ought to ditch Bill's poetry reading and go to the Laker game.

In response to this objection, I am happy to grant that, as far as selfishness goes, Aaron ought to ditch Bill's poetry reading and go to the Laker game, or, we might say, taking only his own interests into account, that is what Aaron ought to do. What this brings out is that "prudence" is not a name for selfishness or for self-interest. Nevertheless, on my view, prudence is indeed distinctly concerned with the agent's own well-being. But prudence is concerned with the responsible and warranted pursuit and protection of one's own well-being.¹⁵ That is, prudence is concerned with the pursuit of one's own well-being that is supported by one's reasons.¹⁶

12 Summary

I've argued that the Common View is false, and I've proposed an alternative view. On My Proposed View, the considerations that make something the morally best thing to do are distinct from the considerations in virtue of which morality *recommends* one action rather than another. Something is a morally good thing to do because of the other-regarding considerations that tell in favor of doing it. Whether something is a morally good thing to do is insensitive to how much it burdens the agent. (Or perhaps the fact that an action burdens an agent more severely

makes it a morally better thing to do!) But what morality *recommends* does take into account burdens on the agent, just as what morality *requires* takes burdens into account. Similarly, on My Proposed View, what prudence *recommends* is sensitive to all the agent's reasons, not just her self-interested reasons.

On My Proposed View, it is sometimes true that, morally, an agent ought to take a particular option, and it is sometimes true that, prudentially, an agent ought to take a particular option. But each of these claims is true only if, all things considered, the agent ought to take that option. On the view I've offered, these claims never conflict: it is never the case that, morally, an agent ought to do something, while prudentially, she ought refrain from doing it. However, it can be true that *both* morally, an agent ought to do something, and also, prudentially, she ought to do the same thing. This will be true whenever moral and prudential considerations both provide central explanations of why an agent ought to do a particular thing.

My Proposed View does not hold that there are three distinct *oughts*, one moral *ought*, one prudential *ought*, and one all-things-considered *ought*. Rather, it is the all-things-considered *ought* that is at play throughout the phenomena we have discussed. Some *ought* facts are *moral* facts in that they are centrally explained by moral considerations.¹⁷ Some *ought* facts are *prudential* facts in that they are centrally explained by prudential considerations.¹⁸

13 Why this matters

Why is it important to see that My Proposed View is true and that both the Common View and the Naïve View are false? Moral philosophers' failure to see the truth of My Proposed View partly explains, and is partly explained by, their being caught in the grip of a false picture of moral requirement and of the nature of the realm of the morally permissible.

This chapter is part of a series of papers¹⁹ in which I have been arguing that philosophers mistakenly endorse the following characterization of moral requirement:

The False Characterization:

An agent is *morally required* to do something just in case, all things considered, she ought to do it, and the reasons that explain why it ought to be done are *moral reasons*.

This characterization is incorrect, however, because of the truth of the following claim (which I introduced in section 8):

The Supererogation-Ought Claim:

Some (but not all) supererogatory options are such that, all things considered, the agent ought to take them. (And some of these are such that the moral considerations that make them morally good to take also centrally explain its being the case that they ought to be taken, all things considered.)

Consider the case of Georgia. In this chapter, I focused on the claim that, morally, Georgia ought to visit her sick neighbor Harriet. But it is also true that, all things considered, Georgia ought to visit Harriet; and moral reasons explain why. Georgia is not morally required to visit Harriet; thus, this case shows the False Characterization to be false.²⁰ If in this case, Georgia fails to visit Harriet, then Georgia makes a *mistake* (she acts as she ought not to act, all things considered); it is a *moral* mistake (in that the reasons she ought not to act this way are moral reasons), and yet it is *morally permissible*. In my recent series of papers, I have been arguing for the importance of recognizing that some actions are *morally permissible moral mistakes*.

Four things are at stake in properly understanding these moral issues. First, it is important that we not endorse the False Characterization of moral requirement; we should not misunderstand the name of moral requirement. Second, it is important that we correctly understand the way that reasons function within the realm of the morally permissible: among morally permissible options, moral reasons continue to have sway, and they can settle that, all things considered, one ought to take a supererogatory option. We must correctly understand the normative status of supererogatory options: some are such that, all things considered, they ought to be taken, but some are not. As I have argued in this chapter, some supererogatory options are such that *morally*, they ought to be taken, but some are not. Third, if we are caught in the grip of the false picture, then we will misunderstand the force of *moral arguments* against behaving in certain ways: we will assume that every moral argument against a way of behaving shows that behavior to be *morally wrong*. But it turns out that a moral argument may merely show a way of behaving to be a moral mistake without showing it to be morally wrong. Fourth and finally, once we abandon the false picture and recognize the existence of *morally permissible moral mistakes*, we can recognize the possibility of new moral views that we have not considered. For example, consider the view that it is a *moral mistake* to eat meat and that eating meat is a morally bad thing to do,²¹ and yet that it is not morally wrong to eat meat. This view might explain what is otherwise puzzling: that some vegetarians refrain from eating meat for moral reasons, and yet they accommodate meat eating in others. It is not in general morally permissible to accommodate others' wrongdoing; but it may be permissible to accommodate others in committing mere moral mistakes. Another example is given by the view that gamete donation (sperm and egg donation) is a moral mistake, though it is also a wonderful thing to do; in other work, I argue that this view is worth taking seriously.²²

We might say that the crucial insight is that the category of *moral mistakes* is bigger than the category of *morally wrong behavior*. Moral mistakes are options such that, all things considered, one ought not to take them, and moral reasons explain why they ought not to be taken; some of these are not morally wrong. But we could equally say that what's crucial is to recognize the *flipside category*: the category of options such that, all things considered, one ought to take them, and moral considerations explain why one ought to take them. This category does not just include morally required options; it also includes some supererogatory options.²³ In this chapter, I've argued that this flipside category is the category of *options that, morally, ought to be taken*.

So now we come to the topic of this chapter. Those who endorse the False Characterization of moral requirement do not see the *flipside category* as a distinct category from the category of morally required options. Thus, when seeking to understand the category of *options that, morally, ought to be taken*, they see only two plausible views: the Naïve View, which equates *options that, morally, ought to be taken* with morally required options, and the Common View, which equates *options that, morally, ought to be taken* with morally best options. It is only in rejecting the False Characterization that one can see a distinct third category: as I've argued in offering My Proposed View, the *options that, morally, ought to be taken* are those options that meet the following condition: all things considered, the agent ought to take the option, and moral considerations explain why it ought to be taken. This includes all morally required options, but it also includes some supererogatory options.²⁴

Notes

- 1 For example, in *Opting for the Best* (New York: Oxford University Press, 2019), Douglas Portmore explicitly develops a view along these lines, arguing that – for maximal (complete) options that one has at a time – morally, one ought to take an option just in case it is one's morally best option, and

prudentially, one ought to take an option just in case it is one's prudentially best option, and Portmore argues that these claims hold, no matter what makes an option morally best (thus, no matter whether morality involves any agent-relative constraints or permissions).

In another example, in his work on the supererogatory, Paul McNamara uses "ought" to pick out an agent's morally best option. (See his "Supererogation: Inside and Out," *Oxford Studies in Normative Ethics* 1 (2011): 202–235, and other work.)

In a third example, Stephen Finlay argues that what one "morally ought to do" is the morally best thing one could do, in his "Too Much Morality," in *Morality and Self-Interest*, ed. Paul Bloomfield (New York: Oxford University Press, 2007), 136–154.

Work on the semantics of the word "ought" often assumes that "ought" picks out an agent's best option; note that my target is the more specific claim that there is a moral *ought* that picks out an agent's *morally best* option.

Note that the idea that, morally, an agent ought to take her morally best option does not commit one to consequentialism: the morally best *option* is simply the morally best *thing to do*, and it need not be the option that has the best consequences. For example, one could hold that one's morally best option is always to keep one's promises, even in cases in which breaking a promise would have better consequences.

- 2 Note that Carl could also express this truth – the truth that, morally, Aaron ought to go to the reading – by simply saying, "You ought to go to the reading."
- 3 Note that Fiona could also express this truth – the truth that, morally, Donna ought to keep Ellen under consideration – by simply saying, "You ought to keep Ellen under consideration."
- 4 At this point in the dialectic, while discussing why we might be drawn to the Common View and should reject the Naïve View, I am simply taking these *ought* claims about Georgia and Isaac to be *true* in these cases. But I should mention that my considered view is *not* that these claims are obviously true given the details of the cases as they've been described. Rather, my considered view is that there are possible versions of these cases (as described) in which these normative claims are true. (That would be enough to imply that the Naïve View is false.)
- 5 An objector might say that James's morally best option is not playing with Kenny but rather doing something else entirely, like driving to the nearest hospital to offer up one of his kidneys to a stranger. My response is that whether an option is a person's morally best option depends on how it compares to its alternatives: the common view holds that, morally, one ought to do something just in case it is morally better than any *alternatives to it*. It's true that donating a kidney is morally better than playing Parcheesi with Kenny, but donating isn't really an alternative to playing with Kenny: James could simultaneously play with Kenny and make a phone call to the hospital to arrange a kidney donation.
- 6 An objector might say that I have misunderstood what it is for an action to be the morally best thing an agent could do: what it is for an action to be morally best *just is* that the action is something that, morally, the agent ought to do. This objector agrees with my claims about whether, morally, an agent ought to do something but disagrees with my claims about which actions are morally best. My reply to this objection is that it divorces whether an action is morally good and morally better from whether an action is praiseworthy or how praiseworthy it is, whereas in fact there is a close connection between these things. For example, it is not the case that, morally, Nicole ought to take the ASL class, but she would be praiseworthy for doing so; indeed, she'd be more praiseworthy than she would be for studying the YouTube videos. Taking the ASL class is a morally better thing to do than studying the YouTube videos, as revealed by the difference in the praiseworthiness of the two actions.
- 7 The existence of these "second-best" supererogatory actions, such as Nicole's watching the YouTube videos, is underappreciated in discussions of the supererogatory. For example, in Terry Horgan and Mark Timmons's, "Untying a Knot from the Inside Out: Reflections on the 'Paradox' of Supererogation," *Social Philosophy and Policy* 27 (2010): 29–63, they gloss the paradox of supererogation as arising because supererogatory actions are "morally best" (29). Similarly, Jamie Dreier's, "Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't," in *Satisficing and Maximizing*, ed. Michael Byron (Cambridge: Cambridge University Press, 2004), 131–145, proposes that supererogatory actions are those that are required from the perspective of beneficence, and Douglas Portmore's, "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11 (2008): 369–388, says "for there is a sense in which supererogatory acts are acts that agents morally ought to perform" (379). Paul McNamara's "Supererogation: Inside and Out: Toward an Adequate Scheme for Common Sense Morality" points out that sometimes a supererogatory action is not an agent's morally best option, that such an action can nevertheless be praiseworthy, and that this phenomenon is underappreciated.

- 8 Consider the following objection: all that's going on in Nicole's case is that, given the expense and time-consumingness of the class, that option has already been ruled out, so it's true that, morally, Nicole ought to study the YouTube videos, simply because we are only evaluating a *proper subset* of her options, and it is the morally best of those options. The problem with this objection is that, if it were correct, then it would *also* be true that, morally, Nicole ought to take the class; this would be true relative to her whole set of options, but it is not true. (The objector is certainly right, of course, that sometimes it is true that, morally, someone should do something, where this is a truth about doing that as compared to a subset of her options. For example, it is true that, morally, a murderer should act quickly (rather than painfully slowly). But in such cases, it is also true that, morally, the murderer should not kill anyone at all.) For a contrastive view of *ought*, see Stephen Finlay and Justin Snedegar's, "One Ought Too Many," *Philosophy and Phenomenological Research* 89 (1) (2014): 102–124. For discussion of the gentle murderer, see Frank Jackson, "On the Semantics and Logic of Obligation," *Mind* 94 (1985): 177–195.
- 9 We can distinguish reasons that are morally relevant from moral reasons. Self-regarding reasons are morally relevant, but they are (at least typically) not moral reasons: they do not tell in favor of an action's being morally good, morally bad, morally required, or morally wrong (thus, they are not moral reasons), though they may tell *against* an action's having one of these properties (thus, they are morally relevant). Douglas Portmore's "Are Moral Reasons Morally Overriding?" points out that we should distinguish reasons that are morally relevant from moral reasons.
- 10 I take this claim about James to follow clearly from the described facts in his case. Laura's case is more subtle; perhaps further details matter, such as how tired she is and how much she has had to deal with sexist comments like these. All I need, to use the case to argue against the Common View, is that there is a version of Laura's case in which it is false that, morally, she ought to kindly explain to Mark why what he said was awful.
- 11 A different objection goes as follows: "It is not obnoxious (or jerky) to say 'morally, you ought to do this' because saying that does not recommend the action; it simply points out something relevant to whether to do it; saying 'morally, you ought to do this' is not the same as saying 'all things considered, you ought to do this.'" This objection misunderstands the argument I make in section 4. My argument is: Only a jerk would recommend taking this option; the recommendations of morality are not recommendations that only a jerk would make; therefore, morality does not recommend taking this option. My argument thus does not rely on the claim that it would be obnoxious (or jerky) to say "morally, you ought to do this" (though I do in fact think that this would be obnoxious, and false, in these cases); rather, my argument relies on the claim that it would be obnoxious to say "I recommend that you do this" or simply to offer the exhortation, "do this."
- This objection brings out something important about the sense in which morality makes recommendations. The recommendation is simply the recommendation *of an action*; the recommendation is not a claim about the moral properties of an action. For example, if it is true that, morally, an agent ought to dance, then morality recommends *dancing*.
- 12 Note that it's clearly true that *if* an agent is morally required to do something, *then* she is such that, morally, she ought to do it. So it's important that My Proposed View accommodate this truth.
- Consideration of this truth can explain why My Proposed View includes the condition that there are *some moral reasons* that *centrally explain* its being the case that one ought to do something, all things considered. Compare My Proposed View to the following view: An agent is such that, morally, she ought to PHI just in case, all things considered, she ought to PHI, and *all the reasons that explain* why she ought to PHI are *moral reasons*. Some morally required actions would not meet this condition, because sometimes when an action is morally required, it also has other decisive reasons in favor of it: an action might be both a morally required promise-keeping but also necessary to save the agent from wasting a great deal of money. In that case, some of the reasons in favor of the action, which explain why it ought to be performed, are not moral reasons. So the comparison view would not imply that, morally, the agent ought to perform this morally required action. By contrast, My Proposed View does accommodate the fact that, morally, the agent ought to perform this morally required action.
- For related discussion, see section VII of my "Morally Permissible Moral Mistakes" (*Ethics*, 2016).
- 13 I argue for the Supererogation-Ought Claim (though not by that name) in my "Morally Permissible Moral Mistakes" (*Ethics*, 2016) and my "Morality Within the Realm of the Morally Permissible" (*Oxford Studies in Metaethics*, 2015).

In "Morally Permissible Moral Mistakes," footnote 18, I briefly considered a view along the lines of My Proposed View in the current chapter, applied to "what one morally ought to do," and I rejected

the view. Though in the current chapter I don't use the locution "what one morally ought to do" (speaking instead of "what, morally, one ought to do"), the view I offer in the current chapter might naturally be applied to that similar locution. Thus, I now think my earlier footnote 18 was mistaken (though the broader point at issue in that footnote was not mistaken – the broader point was that the central lessons of that paper were not already well understood).

14 Susan Wolf, "Moral Saints," *Journal of Philosophy* 79 (1982): 419–439, discusses a "moral saint," someone who always does the morally best thing she could do. Wolf discusses *whether* one should want to be a moral saint; her view is that one should not. I am arguing that not even *morality* urges such a life upon us. Morality urges morally good action only when, all things considered, it is the thing to do: morality sometimes urges us to do the morally best thing we could do, but it often does not.

15 There is some use of "prudent" to mean simply wise or responsible; that is not the use of "prudent" or "prudence" in this chapter.

16 An objection from the opposite direction goes as follows: the Common View would not hold that, prudentially, Professor X ought to sue his students, because it makes one's life go worse when one does something deeply immoral, and it's immoral for him to sue his students. In response to this objection, let's recognize a *narrow* and a *broad* sense of what's good for a person. Suppose it's true that being immoral makes one's life overall worse; that is true in the broad sense of what's good for a person. But we still should acknowledge that sometimes a person does something deeply immoral because it is better for him; this is true in the narrow sense of what's good for a person. The plausible version of the Common View – on which morality and prudence sometimes give conflicting recommendations, and prudence always recommends what's best for the agent – uses this narrow sense of what's good for a person. So the objector is incorrect.

17 An objector might react to my argument as follows:

You're argued that there is no distinctively moral *ought*. But there is a *pro tanto* moral *ought*: whenever someone has a *pro tanto* moral obligation to do something, one can truly say "morally, she ought to do this" even though often it is false that, all things considered, she ought to do it.

Let's grant to the objector that one can truly say "morally, she ought to do this" whenever someone has a *pro tanto* moral obligation. However, the existence of a *pro tanto* moral *ought* would not challenge my argument. In this chapter, I have argued against a moral *ought* such that, morally, an agent ought to something, just in case it is the morally best thing she could do. A *pro tanto* moral *ought* is not the kind of *ought* I am arguing against, because it will often be true that, morally, an agent ought (*pro tanto*) to do something although doing it is not her morally best option.

18 There is a feature of my view that might seem peculiar: from the fact that, morally, S ought to ϕ , and the fact that ψ ing is an alternative to ϕ ing, it does *not* follow that, morally, S ought to refrain from ψ ing. For example, morally, Nicole ought to watch the YouTube videos; taking the ASL class is an alternative to watching the YouTube videos, but it is not the case that, morally, Nicole ought to refrain from taking the ASL class: there is no moral problem with her doing so. But this feature of my view will seem peculiar only if we take there to be a moral *ought* at play in these facts. That is not right. There is simply the all-things-considered *ought* at play in these facts. And it does follow from the fact that, morally, Nicole ought to watch the YouTube videos (and that taking the ASL class is an alternative) that Nicole *ought to refrain* from taking the ASL class: but it is simply true that, all things considered, she ought to refrain from taking the ASL class.

This apparent peculiarity helps to illustrate that on my view, it is really not the case that there is a moral *ought*.

19 See footnotes 13 and 22.

20 Note that I only need the claim that there is some version of Georgia's case in which, all things considered, she ought to visit Harriet. I do not need the claim that the details of the case make it obvious that, all things considered, Georgia ought to visit Harriet. More generally, what I need is the claim that *sometimes* a supererogatory action is the thing to do, the thing that is the correct choice in light of one's reasons: sometimes, all things considered, one ought to perform a supererogatory action.

21 Note that a way of behaving might be a moral mistake without being a morally bad thing to do, so the claim that eating meat is a morally bad thing to do goes beyond the claim that eating meat is a moral mistake. (For example, if Nicole fails to learn any sign language, neither watching the YouTube videos nor taking the class, then she makes a moral mistake in failing to watch the YouTube videos, but she doesn't do anything morally bad. She merely fails to do something morally good that she ought to have done.)

- 22 For extended discussion of these two views, see my “Eating Meat as a Morally Permissible Moral Mistake,” in *Philosophy Comes to Dinner* (New York: Routledge, 2015) and “Gamete Donation as a Laudable Moral Mistake,” in *Oxford Handbook of Population Ethics*, eds. Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns (Oxford: Oxford University Press, Forthcoming).
- 23 If there are any *suberogatory* actions, then this flipside category also includes some (or all) refrainings from suberogatory actions. Supererogatory behavior is what is morally good to do but not morally required; suberogatory behavior is what is morally bad to do but not morally wrong. It is disputed whether any behavior is really suberogatory, even among those who believe in the supererogatory. For further discussion of how the suberogatory fits into this flipside category, see my “Morally Permissible Moral Mistakes,” section V.E. For arguments for and against the suberogatory, see Julia Driver, “The Suberogatory,” *Australasian Journal of Philosophy* 70 (1992): 286–295 and Hallie Liberto, “Denying the Suberogatory,” *Philosophia* 40 (2012): 395–402.
- 24 For helpful comments on drafts of this chapter, I thank Tyler Doggett, Stephen Finlay, Peter Graham, Alex Guerrero, John Horty, Elinor Mason, Douglas Portmore, Mark Schroeder, Ted Sider, Kurt Sylvan, my spring 2017 graduate seminar at Princeton University, and audiences at Brown University; Florida State University; Princeton University; Purdue University; University of Edinburgh; University of Reading; University of Texas, Austin; and University of Southern California.