

Forthcoming in *Philosophical Studies*

Daniel Immerman

Moral Pickles, Moral Dilemmas, and the Obligation Preface Paradox

Abstract. This paper introduces and defends a new position regarding the question of whether it is possible to have conflicting moral obligations. In doing so, it focuses on what I call a moral pickle. By “moral pickle” I mean a set of actions such that you ought to perform each and cannot perform all. Typically, when people discuss conflicting moral obligations, they focus on the notion of a moral dilemma, which is a type of moral pickle involving two conflicting actions. In other words, a moral dilemma is a pair of actions such that you ought to perform each and cannot perform both. As of yet, there is no debate about the possibility of moral pickles over and above the possibility of moral dilemmas. But as I show, there is good reason to think that moral pickles are possible and moral dilemmas are not.

Keywords. Moral pickle; moral dilemma; obligation; belief; preface paradox; obligation preface paradox.

Logan has just been hired in her first philosophy job and, as the sole logic instructor, is responsible for teaching a hundred person logic class every quarter for the foreseeable future. Before she starts grading her first weekly problem set, she asks a mentor: “How carefully should I grade these problem

sets?” The mentor responds: “Each time you grade a problem, you ought to mark it right if it’s right and wrong if it’s wrong, but you don’t have to give comments every time.” Logan responds: “While your advice seems right, you’ve misunderstood me. I’m worried about how carefully I will need to grade so as to avoid making mistakes. After all, if each student turns in ten problem sets a quarter, and each problem set has ten problems, and if I keep teaching this logic class every quarter and stay in this department for the next thirty years of my life, I’m going to be grading a total of about a million problems. I don’t ever want to make even a single mistake in grading. But to avoid doing so, given how many of these problems I’ll have to grade, I’ll probably have to recheck each problem set a number of times.” The mentor responds: “That sounds like a bad idea; the procedure you describe sounds like a massive time commitment. I don’t recommend grading in such a way that you never ever make a mistake. On the contrary, you ought to grade in such a way that you make a mistake now and again – if you graded so carefully that you never ever made a mistake, you wouldn’t leave yourself enough time to fulfill your other obligations.”

Assuming Logan’s mentor’s advice is correct, this case is an example of a moral pickle. By *moral pickle* I mean a set of actions such that you ought to perform each and cannot perform all.¹ In brief, the case of Logan is a case of

¹Here and elsewhere in this paper, in talking about actions one “can” perform or propositions one “can” believe, I mean to be referring to what is logically possible. For example, in the case of Logan, it is logically impossible to both grade each problem correctly

a moral pickle because Logan ought to grade each problem correctly and at the same time ought to grade in such a way that she makes mistakes now and again. Slightly more carefully, the set of actions she's obligated to perform has a million and one members: the first million – one per problem she will grade – are to mark the problem in question right if it's right and wrong if it's wrong, and the last action is to grade in such a way that she makes mistakes now and again.² Logan ought to perform each of these actions but cannot perform all.

and also make mistakes in grading now and again.

²As I interpret the phrase “grade in such a way that she makes mistakes now and again”, Logan will only have graded in this way if she *actually makes* mistakes every now and again. An alternative interpretation is that it merely requires grading in such a way that one is *likely to make* mistakes. I think my interpretation is the more natural one. Some evidence for my interpretation comes from looking at other phrases with a similar structure. For instance, consider the recommendation: “deliver philosophy papers in such a way that the audience laughs from time to time.” On the most natural interpretation of this phrase, in order to deliver one's papers in this way, audience members must actually laugh occasionally during one's philosophy talks. In other words, if I give a number of philosophy talks and no one ever laughs during any of them, I have not delivered my philosophy talks in such a way the audience laughs from time to time. (Just in case the reader is worried that I rest my case for moral pickles on this point about interpretation, I should note that I offer arguments later on in this paper for the claim that Logan ought to make a mistake every now and again, as opposed to merely having the obligation to grade in such a way that she's likely to make a mistake every now and again.) Thanks to an anonymous referee for pressing me on this.

I should note that by “moral” in “moral pickle” I mean to emphasize that I’m talking about moral obligations, as opposed to other sorts of obligations, such as practical obligations.³

In addition, I should note that I mean to be using the all-things-considered moral ought. By “all-things-considered”, I mean to contrast my sense of “ought” with a “prima facie” or “pro tanto” sense of “ought” – my sense of “ought” is not merely one that applies at first glance, or one that can be overridden by other obligations; rather, it is a sense that survives further scrutiny and is not overridden.

The notion of moral pickle should be contrasted with the notion of a moral dilemma. “Moral dilemmas”, as I shall use the term, are a species of moral pickle involving only two obligations. In other words, a *moral dilemma* is a set of two actions such that you ought to perform each and cannot perform both.⁴ It is currently deeply contentious whether moral dilemmas

³I am assuming that, for each problem, Logan has a *moral* obligation to grade it correctly. One reason to think this is because grading it incorrectly would be unfair to the student who submitted it. If you nonetheless doubt that this is a moral obligation, feel free to alter the example to make it more obviously a moral one. One can substitute in any minor, easy-to-perform task, for which it’s possible that someone will make a mistake due to lack of attention. For instance: turning off light switches to save electricity, transcribing credit card numbers for people who have decided to donate to some charity, etc.

⁴There are other interesting distinctions in the neighborhood. For example, one can distinguish between “higher-order moral pickles” and “lower-order moral pickles”. To motivate the distinction, return to the example in which Logan is told she should grade in

are possible.⁵

As of now, there isn't a debate about the possibility of moral pickles over and above the possibility of moral dilemmas; indeed, as of the writing of this paper, there wasn't even a name for moral pickles. Sometimes people writing on moral dilemmas fail to carefully distinguish between moral pickles and moral dilemmas. For example, in Earl Conee's classic article "Against Moral Dilemmas," Conee starts the paper by defining the expression "moral dilemma" in such a way that it is equivalent to what I have been calling

such a way that she make mistakes now and again. This obligation is (arguably) "higher-order" where a higher-order obligation is an obligation to act in such a way as to bring about other actions. In particular, the action in question is presumably the adoption of some sort of procedure that will in turn lead her to grade problems in a certain way. A higher-order moral pickle, then, is a moral pickle in which some of the obligations are higher-order, whereas a first-order moral pickle is one where none of the actions are.

The obligation to make mistakes is also (arguably) "complex" in that it involves multiple actions. An obligation is "simple" if it involves just one action. (There are presumably multiple ways to individuate what counts as a single action, but on at least some of these, grading in such a way that one makes mistakes from time to time counts as multiple actions, one for each problem one grades.) This yields another distinction: a "complex moral pickle" is one on which some of the actions are complex, and a "simple moral pickle" would be one on which all the actions are simple. Thanks to an anonymous referee for encouraging me to discuss these related distinctions.

⁵This debate seems to have been kicked off by Bernard Williams in [Williams, 1965]; key works include [Conee, 1982] and [Marcus, 1980]. For a good list of sources on the debate, see the bibliography of [McConnell, 2014].

“moral pickles”, writing, “Call an agent’s predicament a ‘moral dilemma’ just when the agent cannot do everything that is morally obligatory for him to do in the situation, though he can carry out each obligation” [Conee, 1982, 87]. A page later, he uses the same expression to start talking about what I have been calling “moral dilemmas”, writing, “If there are moral dilemmas, then some truths have this form: $O(A) \wedge O(\neg A)$ ” [Conee, 1982, 88].⁶ But I do think it is important to distinguish between the two. In particular, as I shall argue, there are good reasons to think that moral pickles are possible and moral dilemmas are not.

In the course of arguing for this conclusion, I shall draw on some parallels between moral obligation and belief. This is because the view I focus on – that moral pickles are possible and moral dilemmas are not – is analogous to a view in epistemology that is widely held. Furthermore, certain key arguments for this epistemic view carry over to the ethical domain, as we

⁶Given that attention has not been paid to the distinction between moral dilemmas and moral pickles, there has been some variation in the way “moral dilemma” has been defined in the literature. Often, it is defined as I have defined it (see e.g. [Brink, 1994, 214-5], [Haan, 2001, 269], [Marcus, 1980, 122], [Rajczi, 2002, 310], [Sinnott-Armstrong, 1987, 128]), but sometimes people use the word “moral dilemma” as I use the word “moral pickle” (see e.g. [McConnell, 2014].). I don’t wish to quibble over which words one uses to pick out which things, as long as these things are properly distinguished. But perhaps I should offer a modest defense of my choice of terminology: in addition to capturing the way that most have used the term “moral dilemma”, it also has etymological advantages, as “di” means two.

shall soon see.

Let me briefly preview. The epistemic analog of the view that moral pickles are possible is that sometimes one ought to have inconsistent beliefs. In other words, the epistemic analog says that sometimes it is the case that one ought to believe each of a set of propositions even if it's impossible for them all to be true. This view is widely held. In addition, one of the major arguments for it, which relies on the so-called "preface paradox", carries over to the ethical domain, as I shall show in the first section of my paper. Meanwhile, the epistemic analog of the view that moral dilemmas are impossible is that one never ought hold a pair of inconsistent beliefs. In other words, the epistemic analog says it is never the case that one ought to believe two propositions that are inconsistent. And again, some of the reasons for holding this view carry over to the moral domain, as I shall show in the second section of my paper.

Section 1. How the obligation preface paradox supports the possibility of moral pickles

In this section, I will show that there is good reason to think that moral pickles are possible. As I note above, my argument is closely connected to the preface paradox.⁷ The preface paradox arises because there seem to be

⁷The preface paradox was originally introduced by D.C. Makinson [Makinson, 1965].

cases in which you ought to believe each of a number of propositions and at the same time ought to believe that at least one of them is false. I will be discussing a closely-related paradox that I will call the “obligation preface paradox.” It arises because there seem to be cases in which you ought to perform each of a number of actions and also ought to perform an action that will keep some of these other actions from being performed.⁸

Section 1.1 Introducing the preface paradox

Let me start by stating the preface paradox. Knowing about the preface paradox will be useful because certain key arguments regarding the preface paradox carry over to the obligation preface paradox and help lend support

For a good list of sources on the debate, see the bibliography of [Sorensen, 2011].

⁸So far as I know, no one has offered a version of the preface paradox for moral obligation. That said, recent papers by Simon Goldstein [Goldstein, 2016] and Sam Shpall [Shpall, 2016] have focused on similar issues; they both offer a version of the preface paradox for intention. It is perhaps worth briefly noting some ways in which our papers differ. In particular, they focus on whether one can be *rationally* obligated to hold conflicting *intentions*, while I am focused on whether one can be *morally* obligated to perform conflicting sets of *actions*. This difference in focus yields important differences regarding the upshots of our papers; the key upshot of my paper is that the literature on moral dilemmas has neglected an important but attractive alternative, namely the alternative on which moral dilemmas are impossible but moral pickles are possible.

to the claim that moral pickles are possible.

The preface paradox involves a story about a careful historian who writes a book with a large number of claims, each of which is carefully researched. Even though the book is carefully researched, it's likely that it contains at least one error – even the most carefully researched books of this sort do. So the historian notes in her preface that her book contains at least one mistake and says that the fault is entirely her own.

The reason this is a paradox is because each of the following claims seems to be true, and yet they are mutually inconsistent:

1. For each proposition in her book, the historian ought to believe it. (After all, the book is carefully researched.)
2. The historian ought to believe that at least one of these propositions is false. (After all, there are many claims in the book and thus it is extremely likely that one is false.)
3. It is impossible for there to be a person and a set of propositions such that the propositions can't all to be true at once and nonetheless the person ought to believe each of them.

A popular response to the preface paradox is to reject the third claim.⁹ The reason this strategy is popular is because rejecting either of the other two claims comes with significant costs; I will discuss some of these costs momentarily.

⁹That many respond in this way is noted in e.g. [Sorensen, 2011].

Section 1.2 Introducing the obligation preface paradox and defending its first claim

Now that I've introduced the preface paradox, I'm ready to introduce the obligation preface paradox. The obligation preface paradox involves a situation in which a subject ought to perform each of a set of actions and at the same time ought to perform an action that will keep some of these other actions from being performed. To develop it, I will use the example from my opening paragraphs involving Logan.

As with the preface paradox, there are three claims that seem to be true and yet are mutually inconsistent. I will explain and motivate the first two and then state the final one. Here is the first claim:

1*. For any given problem she's grading, Logan ought to mark it right if it's right and wrong if it's wrong.

This claim is parallel to claim 1 from the preface paradox, the claim that ran "For each proposition in her book, the historian ought to believe it." I mentioned above that rejecting claim 1 comes with significant costs. And the same thing is true of 1*; I shall now describe three such costs.

First, there is an intuitive cost in reject 1*. After all, as the case is set up, the problems are relatively simple to grade. In addition, it will be helpful

to the students to mark the problems correctly, and there's nothing to be gained by marking them incorrectly. In such a case, it is intuitive that she ought to mark them right if they're right and wrong if they're wrong.¹⁰

The second cost to rejecting claim 1* is that rejecting it would seemingly lead to a widespread shift in our practices regarding obligation and the norms that guide them. There are a couple of related reasons to think that this is so.

One is that it there seems to be fairly strong reason to perform the actions described in 1*, that is, fairly strong reason to mark these problems right if they're right and wrong if they're wrong. After all, this is relatively easy to do, it benefits the students, and there's nothing to be gained from marking them incorrectly. If you say that it's not the case that Logan ought to perform these actions – even though she has such good reason to perform them – it's hard to see how you can avoid the conclusion that there are many actions that we currently think we ought to perform such that it's not the case that we ought to perform them.

Another reason rejecting claim 1* would lead to a widespread shift is that the obligation preface paradox can be generated with other sets of obligations. There is nothing special about grading; so long as we take a sufficiently large number of your obligations, there is good reason to think that it's not the case that you ought to perform them all. (I develop this point in more detail

¹⁰For those who develop parallel points about rejecting 1, see [Christensen, 2004, 34], [Foley, 2009, 37].

in Section 1.5.)

In short, rejecting 1* seems to suggest a rather widespread revision of our obligations and the norms that guide them. And this seems like a significant cost; we should approach such a revision with trepidation.¹¹

To understand the third cost, let us start by looking at a problem with rejecting claim 1. If one were to reject 1 and say that it's not the case that the historian ought to believe the propositions in her book, then one is faced with the question: what attitude ought she take towards these propositions? And the options do not look particularly good. For instance, one might propose that it's not the case that the historian ought to believe *outright* the claims in her book, but instead ought to hold beliefs concerning their *likelihoods*. But such a proposal faces a number of difficulties. This is partially true for practical reasons; it's more difficult to remember and reason about beliefs concerning likelihoods than outright beliefs, and thus would be far more difficult to figure out what we ought to believe if we made such a switch. In addition, beliefs are integrated with others of our states, such as knowledge, and it is not clear how to integrate such states if we switch to recommending beliefs concerning likelihoods.¹²

Likewise, if one were to reject 1* and say that it's not the case that Logan

¹¹For those who develop parallel points about 1, see e.g. [Foley, 2009, 44], [Worsnip, 2016].

¹² For those who make these points, see e.g. [Christensen, 2004, 12-3], [Foley, 2009, 45-6], [Harman, 1986, 26-7].

ought to mark a problem right if it's right and wrong if it's wrong, then one is faced with the question: what ought she do regarding these problems? And the options do not look particularly good. For instance, one could say that for each problem, Logan ought to grade the problem in such a way that she's likely to mark it right if it's right and wrong if it's wrong. But as with the belief claim, such a proposal faces a number of difficulties. This is partially true for practical reasons, it's more difficult to remember and reason about obligations to make outcomes more likely than it is to reason about obligations to bring about outcomes, and thus far more difficult to figure out what we ought to do if we made such a switch. In addition, obligations are integrated with others of our states, such as intention, and it is not clear how to integrate such states if we switch to recommending obligations to make certain things likely.

In sum, as we have seen, there are significant costs to rejecting 1*.

Section 1.3 Defending the second claim of the obligation preface paradox

The next claim in the obligation preface paradox is the following:

2*. Logan ought to grade in such a way that she makes a mistake every now and again.

This claim is parallel to claim 2, the claim that ran “The historian ought to believe that at least one of these propositions is false.”

As with 1*, there are significant costs to rejecting claim 2*. First, claim 2* is intuitively plausible.¹³ After all, if Logan grades in such a way that she never makes a mistake, she’ll spend enormous amounts of time grading, time that would be better spend elsewhere. And thus, it seems intuitive plausible to say that Logan ought to grade in such a way that she makes mistakes now and again.

To further develop the point about intuitive plausibility, it’s worth noting that we can naturally fill in the story in such a way that the mentor is explicit that the obligation is to grade imperfectly. For example, we can imagine Logan saying to the mentor: “Ok, I will grade somewhat quickly. But how quickly should I grade? Should I grade so quickly that I make mistakes on nearly every problem set? Or should I grade more carefully than that?” And to continue the story we can imagine the mentor replying: “That is too many mistakes! You should make a mistake every twenty problem sets or so. If a semester goes by in which you don’t make any mistakes, you’re grading more carefully than you should and you should speed up; if you start making mistakes on every other problem set, you’re grading too carelessly and should slow down”.

In addition, as with claim 1*, rejecting claim 2* and replacing it with

¹³For those who note a parallel claim about 2, see e.g. [Christensen, 2004, 41], [Makinson, 1965, 205].

another claim leads to problems. To understand them, it helps to understand parallel problems that arise regarding the strategy of rejecting claim 2 and replacing it with another.

In particular, someone might propose that it's not the case that the historian ought to *outright believe* that she made a mistake, but simply that the historian ought to believe that it's *highly likely* she made a mistake.

But if one takes this path, one faces a dilemma, which concerns the question: ought the historian believe that she didn't make a mistake?¹⁴

One option is to say no. On this option, the historian doesn't have an obligation in either direction; it's not the case that she ought to believe the book contains errors and it's not the case that she ought to believe the book doesn't contain errors. Taking this option causes a problem regarding a principle called Belief Agglomeration, which runs as follows:

Belief Agglomeration: If one ought to believe A and one ought to believe B then one ought to believe A and B.

This principle is – according to some – quite plausible. And the main motivation people have when they reject claim 2 is to avoid having to reject Belief Agglomeration. But if you say no to the dilemma, then it is extremely difficult to avoid rejecting Belief Agglomeration. In particular, assuming claim 1 is right, and the historian ought to believe each of the individual propositions

¹⁴In my discussion of the dilemma, I am repeating points made by David Christensen; see [Christensen, 2004, 33-68].

in her book, it follows from Belief Agglomeration that she ought to believe the conjunction and if she believes this then presumably she also ought to believe that she has not made a mistake. (I say assuming claim 1 is right; I've already outlined the costs that come from rejecting claim 1.) In short, if you take the first horn of the dilemma, it is extremely difficult to avoid rejecting Belief Agglomeration.

Let me compare this option with the option of accepting claim 2. If you accept claim 2 (and claim 1) then you also have to reject Belief Agglomeration for the same reason as if you reject claim 2 and say no to the dilemma. But the difference is that at least you get to accept claim 2 and thus avoid the counter-intuitive costs of denying it.

To put the point slightly differently: if one chooses to reject claim 2, this is counter-intuitive, but at least carries with it the promise of being able to keep Belief Agglomeration. But if one says no to the dilemma, one has to reject Belief Agglomeration anyway, thus taking a second hit. There seems little to recommend this position over just affirming claim 2.

The second horn of the dilemma is to say that the historian ought to believe that her book contains no errors. But this horn is, if anything, even worse. For starters, it is deeply counter-intuitive that the historian ought to believe her book contains no errors. It seems extremely arrogant and bizarre for the historian to believe she made not a single mistake.

Furthermore, we were considering a strategy in which one rejects claim 2 and replaces it with the claim that the historian ought to believe that it's

highly likely that her book contains errors. So the option we're considering involves saying both that the historian should believe it's highly likely that her book contains errors and that the historian should believe her book contains no errors. And this seems, if anything, even more deeply counter-intuitive.

In sum, if one rejects 2 and replaces it with the claim that the historian ought to believe it's highly likely that her book contains an error, one faces a dilemma and neither option looks promising.

Similar things hold for 2*. One can reject 2* and retreat to the claim that Logan ought to grade in such a way that she's *highly likely* to make a mistake now and again. But then one faces a dilemma: is it the case that Logan ought to grade in such a way that she never makes a mistake or not?

One option is to say that it's not the case that she ought to grade so that she never makes a mistake. Taking this option causes a problem regarding a principle called Obligation Agglomeration, which runs as follows:

Obligation Agglomeration: If one ought to perform A and one ought to perform B then one ought to perform A and B.

As with the belief case, in saying no to the question posed by the dilemma, one puts oneself in an uncomfortable middle ground. One has already rejected claim 2*, even though it was intuitive. But at least, so it seemed, one could maintain Obligation Agglomeration. But now one has gone ahead and rejected Obligation Agglomeration as well.

The second horn of the dilemma is to say that Logan ought to grade in such a way that she never makes a mistake. But, again, this is highly counterintuitive. And again, it seems even more deeply counter-intuitive to say both this and say that she ought to grade in such a way that she's highly likely to make a mistake from time to time.

Let me complete my defense of 2* by raising and then responding to an argument against it. This argument runs as follows: if Logan grades the first 999,999 problems correctly, it's not the case that she's suddenly obligated to grade the last one incorrectly. And this shows that she's not obligated to make a mistake every now and again.¹⁵

In response, I would like to challenge the implicit premise of this argument, viz. that if one ought to do something, then in every possible situation, no matter how unlikely, in which one starts doing it and can finish doing it, one ought to finish doing it.

This implicit premise seems to generate highly counterintuitive results. For example, suppose I ought to feed my two cats their evening meal. I feed one of them but run out of food before I can feed the other. So I go to the pet store and buy cat food, then come home. I know that my one cat hasn't been fed, but I don't bother feeding it and it goes hungry as a result. Suppose, on discovering that I have not fed the one cat, my partner is angry and complains "you ought to have fed both cats." Suppose I then respond: "actually, I in fact had no obligation to feed both cats. After all, while I was

¹⁵Thanks to an anonymous referee for raising this.

gone, a robber could have broken in and fed the other cat. And if that had happened, I would have had no obligation to feed it. Of course,” I continue, “as a matter of fact no robber broke in. But the fact that one could have broken in and fed the cat shows that I lacked the obligation to feed both cats.” This seems like a unconvincing response. I had the obligation to feed both cats, and failed in my obligation to do so.¹⁶

In other words, I had an obligation to feed both cats even though there was a possible situation – the robber breaking in and feeding my second cat while I was gone – in which I started the process of feeding both cats, could have finished it, and would not have been obligated to finish it. So the implicit premise that generates the argument against 2* is false.

It’s perhaps worth noting that similar things apply in the belief case. Suppose that we affirm 2, that is, we affirm that the historian ought to believe her book contains an error. As the historian is about to form this belief, she starts to imagine the following possibility. Suppose, in the future, some expert committee embarked on a years-long project of examining the

¹⁶Perhaps one might respond that my obligation was to ensure that both cats were fed and that this obligation would hold even if the robber broke in (though of course it would be easier to fulfill in that case). But this is subject to the same sort of worry. It was possible for one of my cats to die, in which case there would be no need to feed it. Maybe the obligation was to ensure that both of my cats were fed, so long as they were still alive? But what if one of them suddenly developed some condition that meant if it ate, it would get deathly ill? And so on. Far simpler and more natural just to say that, because none of these unlikely events took place, I had an obligation to feed both cats.

claims in her book, starting with the first and moving to the last. Suppose they are extremely scrupulous and thus if they affirm one of the historian's claims, it's true beyond a shadow of a doubt. Suppose that they have just finished considering the second-to-last claim and have yet to find an error.

At that point, it would be bizarre for the historian to think: "my book contains a mistake somewhere and the committee has yet to discover it, so the last claim in my book must be false." But the fact that this conditional claim is true – the fact that it is true that were a committee to examine the historian's book and make it to the last claim without having yet found an error, then the historian would no longer have an obligation to believe her book contained an error – this fact does nothing to show that right now – even before any committee has started any process – the historian ought not believe her book contains no error.

Finally, a brief note about the connection between 2* and moral pickles. In order for my overall argument to be successful in demonstrating the existence of moral pickles, I need 2* to describe an obligation to perform an action. But someone might complain that 2* does not describe an obligation to perform an action, rather it describes an obligation to perform an action *in a certain way*. In particular, it describes an obligation to perform an action – grading – in such a way that one makes a mistake every now and again.

A related complaint is that my argument describes a "complex action" – an action that involves a number of simple actions – grading individual problems – as components. And someone might doubt whether one can have

obligations regarding complex actions.¹⁷

In response, it's worth noting that it is possible to alter the example so that 2* more clearly describes an obligation to perform a simple action. For instance, suppose the only way Logan will be able to avoid compulsively checking the problem sets, but instead grade them quickly enough that she makes mistakes from time to time, is to take some a certain sort of pill. Then she ought to take the pill.

Section 1.4 The last claim of the obligation preface paradox

Here is the last claim of the obligation preface paradox:

3*. It is impossible for there to be a person and some set of actions such that the person can't perform all the actions and nonetheless the person ought to perform each.

I have just argued that 1* and 2* are plausible, and from 1* and 2* being true, it follows that 3* false. But if 3* is false, then moral pickles are possible; 3* is just the denial of the claim that moral pickles are possible.

¹⁷Thanks to an anonymous referee for pressing this worry.

Section 1.5 Moral pickles are widespread and of a theoretically interesting sort

It is worth briefly noting that there is good reason to think that not only are moral pickles possible, they're widespread and of a theoretically interesting sort.

First, regarding the claim that they're widespread: while my example from my official statement of the paradox is somewhat unusual, in that not everyone has to grade so many logic assignments, there are analogous examples that apply widely.

After all, we may not all be logic teachers, but we all have numerous obligations. And just as Logan ought to grade in such a way that she makes a mistake from time to time, we ought to act in such a way that we fail to fulfill our obligations from time to time. Acting in such a way that we fulfilled all our obligations would require constantly checking and rechecking that we had successfully fulfilled each obligation, which would be an enormous time sink and prevent us from doing other, more important things. More generally, given that perfection is not always possible with regards to our first-order obligations, it is often a good idea for us to put a procedure in place that will allow us to fulfill our first-order obligations as best as possible. These properties – lack of perfection and need for a procedure – hold not only for those grading logic problem sets but also doctors seeing too many patients, public defenders having too many clients, social workers who have

been assigned too many cases, and so on.¹⁸

Next, I will argue that the moral pickle I've provided in my argument is of a theoretically interesting sort. A typical response to arguments for the possibility of moral dilemmas is to grant that there can be moral dilemmas of a certain sort, but to deny that there can be moral dilemmas of another sort. I will now look at two such distinctions as they would apply to moral pickles and show that my argument shows that there can be moral pickles of the sort deemed impossible.

First, some make a distinction between epistemic and ontological conflicts. Epistemic conflicts hold when an agent does not know which of the conflicting requirements takes precedence, while ontological conflicts hold when there's no fact of the matter as to which takes precedence. Some theorists are happy to allow for moral dilemmas involving epistemic conflicts, but not moral dilemmas involving ontological conflicts [McConnell, 2014].

My example is one of ontological conflict – even with full information about the scenario, the obligations remain. For instance, even if upon Logan's retirement, the mentor went through every problem set Logan had graded and discovered every single error, the mentor's advice would still seem apt; it would still seem as if, for each problem, Logan ought to have marked it right if it was right and wrong if it was wrong and that she ought to have graded in such a way that she made a mistake every now and again.

¹⁸Thanks to an anonymous referee for suggesting that I expand this paragraph and offering some advice on how to do so.

Second, some make a distinction between self-imposed conflicts and conflicts imposed by the world. Self-imposed conflicts arise because of the agent's own wrongdoing while conflicts imposed by the world do not. These theorists are happy to allow for self-imposed moral conflicts but not conflicts imposed by the world [McConnell, 2014].

My example is a pickle imposed by the world: Logan is not in the situation of having to grade logic problems because of past wrongdoing. (Of course, Logan does possess some imperfections that keep her from grading perfectly without a lot of effort. But agents in moral dilemmas are always somewhat imperfect. For instance, a classic example of a moral dilemma involves a character named Sophie who is given a choice by the Nazis: she must choose one of her two children to save and the other will be killed. This dilemma only arises because Sophie is somewhat imperfect; she cannot save both her children because she can't single-handedly fight off the Nazis.)

In sum, there is good reason to think that moral pickles are widespread and of a theoretically interesting sort.

Section 2. Even given my argument, there is still good reason to hold that moral dilemmas are impossible

Suppose one thinks that moral pickles are possible. What positions can one hold regarding the possibility of moral dilemmas?

One option is clear: it is definitely possible to hold that moral dilemmas are possible. After all, moral dilemmas are a type of moral pickle, so everyone who thinks moral dilemmas are possible will automatically think moral pickles are possible.

In this section, I will be arguing that this isn't one's only option. There is good reason to think that moral pickles are possible and moral dilemmas are not. In particular, I will now show that my defense of the possibility of moral pickles is consistent with the premises of a key argument that moral dilemmas are impossible. This argument against the possibility of moral dilemmas makes use of the following two principles:¹⁹

Deontic Consistency:²⁰ If an action is obligatory then it's not the case that you ought not perform it.

Deontic Logic:²¹ If in performing an action you'll bring about a consequence, then if performing the action is obligatory, bringing about the consequence is too.

These two principles entail that there cannot be moral dilemmas.²² But they

¹⁹This argument is presented in e.g. [Brink, 1994, 228], [McConnell, 2014].

²⁰Also known as the weak obligation principle [Horty, 2012, 93].

²¹Also known as closure under logical consequence [Horty, 2012, 83].

²²Here's the argument: Suppose there is a moral dilemma, i.e. two actions A and B such that one ought to perform each but cannot perform both. Then because one cannot perform both, performing A entails not performing B, so by Deontic Logic, if A is

are consistent with the claim that moral pickles are possible.²³ Thus, so long as there is good reason to think both that my argument regarding moral pickles is successful and that these principles are true, there is good reason to think that moral pickles are possible and that moral dilemmas are not.

That said, it is worth briefly noting that my claim that moral pickles are possible is not consistent with the premises of every argument that moral dilemmas are impossible. For example, one important argument against the possibility of moral dilemmas invokes the following two principles:²⁴

Obligation Agglomeration: If one ought to perform A and one ought to perform B then one ought to perform A and B.

Ought-implies-can: If one ought to perform A then one can perform A.

In endorsing the possibility of moral pickles, I have to reject one of these
obligatory, then so is not performing B. But because we're assuming that performing A is obligatory, it follows that not performing B is obligatory. But we have assumed Deontic Consistency and that performing B is obligatory so it follows that it's not the case that not performing B is obligatory. But now we have a contradiction; not performing B can't both be obligatory and not obligatory.

²³For instance, suppose there are three actions, A, B, and C. Suppose that one ought to perform each, but it's not the case that one ought to perform any pair, or all three. Suppose that one can perform any two but can't perform all three. Then both Deontic Consistency and Deontic Logic hold, as can be easily checked.

²⁴For discussion of this argument, see e.g. [Brink, 1994, 228], [McConnell, 2014], [Williams, 1965].

two principles. Assuming my argument for the possibility of moral pickles is sound, and that the argument against moral dilemmas invoking Deontic Consistency and Deontic Logic is sound, it will follow that Obligation Agglomeration is false. This is because, if Obligation Agglomeration were true, then one could, by repeated applications of Obligation Agglomeration transform any moral pickle into a moral dilemma. The way to do this would be to combine the particular obligations from the moral pickle, one by one, into one big obligation.

Obligation Agglomeration is quite controversial [McConnell, 2014]. One interesting result of my paper is that it yields a new argument against Obligation Agglomeration, viz. combining my argument for the possibility of moral pickles with the argument against moral dilemmas invoking Deontic Consistency and Deontic Logic.

Of course, this argument will only be successful if there is good reason to think Deontic Consistency and Deontic Logic are more plausible than Obligation Agglomeration. But I think so and I am not the only one – Terrence McConnell [McConnell, 2014] and David Brink [Brink, 1994, 228] agree.

The idea that moral pickles are possible even if moral dilemmas are not is buttressed by an analogy with beliefs. As I've indicated, thanks to the preface paradox many think that there can be cases in which one ought to hold each of a set of inconsistent beliefs. Far fewer say that there can be cases in which one ought to hold each of a pair of inconsistent beliefs. If

these people wish to endorse analogous claims regarding obligations, they should hold that moral pickles are possible but moral dilemmas not.

The analogies run deeper. Many endorse these claims about beliefs because they hold the view that one ought to believe each of a set of propositions only if each is more likely than not to be true. This means it will never be the case that one ought to believe two things that cannot both be true. For instance, it's not the case, after a coin is flipped, that you ought to believe that it landed heads and you ought to believe that it landed tails, because these are not both more likely than not. But it allows that one can have the obligation to believe multiple things that cannot all be true. For instance, for all I've said about likelihoods, perhaps, after a million-sided die is rolled, you ought to believe it didn't come up one, and you ought to believe it didn't come up two, and so on. After all, these propositions – it didn't come up one, it didn't come up two, and so on – are each more likely than not.

Likewise, one might think that one ought perform each of a set of actions only if the following holds: there is some way of acting that makes it more likely than not that one will successfully perform each action.²⁵ This means that it will not be the case that there are two actions such that one ought to perform each but cannot perform both. But it allows that there can be multiple actions such that one ought to perform each and cannot perform

²⁵For discussion, see e.g. [Brennan and Southwood, 2007], [Estlund, 2011, 212], [Lawford-Smith, 2013], [Southwood, 2015, 519 n. 220], [Southwood and Wiens, 2016], [Southwood, 2016, 11-3], [Vihvelin, 2004, 437-40].

all.

This sort of view helps give a principled reason for rejecting Obligation Agglomeration. Recall that Obligation Agglomeration says that if one ought to perform A and one ought to perform B, one ought to perform A and B. But if we endorse these ideas connecting obligations and likelihoods, we can see that there might be examples in which one can act in such a way that it is more likely than not that one will perform A, more likely than not that one will perform B, but not more likely than not that one will perform A and B.

In sum, as I have shown in this paper, there is good reason to think that moral pickles are possible and moral dilemmas are not. To do this, I first offered an argument that moral pickles are possible and furthermore that they are widespread and of a theoretically significant sort. Next, I argued that even if one accepts all of this, one still has good reason available for thinking that moral dilemmas are impossible.

Acknowledgements

Thanks for helpful comments to Earl Conee, Ryan Hammond, Ting Lau, Caleb Perl, Will Smith, Fritz Warfield, several anonymous referees, and an audience at the 2016 Pacific APA.

References

- [Brennan and Southwood, 2007] Brennan, G. and Southwood, N. (2007). Feasibility in action and attitude. *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*.
- [Brink, 1994] Brink, D. (1994). Moral conflict and its structure. *The Philosophical Review*, 103:215–47.
- [Christensen, 2004] Christensen, D. (2004). *Putting Logic in its Place; Formal Constraints on Rational Belief*. Oxford University Press.
- [Conee, 1982] Conee, E. (1982). Against moral dilemmas. *The Philosophical Review*, 91(1):87–97.
- [Estlund, 2011] Estlund, D. (2011). Human nature and the limits (if any) of political philosophy. *Philosophy and Public Affairs*, 39:207–37.
- [Foley, 2009] Foley, R. (2009). Beliefs, degrees of belief, and the lockean thesis. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*. Springer.
- [Goldstein, 2016] Goldstein, S. (2016). A preface paradox for intention. *Philosophers' Imprint*, 16(14):1–20.
- [Haan, 2001] Haan, J. D. (2001). The definition of moral dilemmas: A logical problem. *Ethical Theory and Moral Practice*, 4(3):267–84.
- [Harman, 1986] Harman, G. (1986). *Change in View*. MIT Press.
- [Horty, 2012] Horty, J. F. (2012). *Reasons as defaults*. Oxford University Press.
- [Lawford-Smith, 2013] Lawford-Smith, H. (2013). Understanding political feasibility. *Journal of Political Philosophy*, 21:243–59.
- [Makinson, 1965] Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25:205–7.

- [Marcus, 1980] Marcus, R. B. (1980). Moral dilemmas and consistency. *The Journal of Philosophy*, 77(3):121–36.
- [McConnell, 2014] McConnell, T. (2014). Moral dilemmas. *Stanford Encyclopedia of Philosophy*.
- [Rajczi, 2002] Rajczi, A. (2002). When can one requirement override another? *Philosophical Studies*, 108(3):309–26.
- [Shpall, 2016] Shpall, S. (2016). The calendar paradox. *Philosophical Studies*, 173(3):801–25.
- [Sinnott-Armstrong, 1987] Sinnott-Armstrong, W. (1987). Moral dilemmas and ‘ought and ought not’. *Canadian Journal of Philosophy*, 17(1):127–40.
- [Sorensen, 2011] Sorensen, R. (2011). Epistemic paradoxes. *Stanford Encyclopedia of Philosophy*.
- [Southwood, 2015] Southwood, N. (2015). Democracy as a modally demanding value. *Noûs*, 49:504–21.
- [Southwood, 2016] Southwood, N. (2016). Does “ought” imply “feasible”? *Philosophy and Public Affairs*, 44(1):7–45.
- [Southwood and Wiens, 2016] Southwood, N. and Wiens, D. (2016). “actual” does not imply “feasible”. *Philosophical Studies*, 173(11):3037–3060.
- [Vihvelin, 2004] Vihvelin, K. (2004). Free will demystified: A dispositional account. *Philosophical Topics*, 32:427–50.
- [Williams, 1965] Williams, B. (1965). Ethical consistency. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 39:103–24.
- [Worsnip, 2016] Worsnip, A. (2016). Belief, credence, and the preface paradox. *Australasian Journal of Philosophy*, 94(3):549–62.