

Maximalism vs. Omnism about Reasons*

Douglas W. Portmore

Abstract: The performance of one option can entail the performance of another. For instance, I have the option of baking a pumpkin pie as well as the option of baking a pie, and the former entails the latter. Now, suppose that I have both reason to bake a pie and reason to bake a pumpkin pie. This raises the question: Which, if either, is more fundamental than the other? Do I have reason to bake a pie because I have reason to perform some instance of pie-baking—perhaps, pumpkin-pie baking? Or do I have reason to bake a pumpkin pie because I have reason to bake a pie? Or are they equally fundamental, as they would be if, say, I had reason to do each because each would have optimal consequences? The aim of this paper is to compare two possible answers to this question—omnism and maximalism—and to argue that the latter is preferable. Roughly speaking, maximalism is the view that only those options that are not entailed by any other option are to be assessed in terms of whether they have some feature (such as that of having optimal consequences), whereas omnism is the view that all options are to be assessed in terms of whether they have this feature. I argue that there are at least two reasons to prefer maximalism, for it is able to overcome two critical problems with omnism.

Keywords: closure principles, deontic logic, maximalism, morality, omnism, options, performance entailment, reasons, rationality, transitivity.

The performance of one option can entail the performance of another. That is, the performance of the one can logically necessitate the performance of the other. For instance, I have both the option of baking a pie and the option of baking a pumpkin pie, and baking a pumpkin pie entails baking a pie, for the statement ‘I bake a pumpkin pie’ entails ‘I bake a pie’.¹ Such instances of performance entailment are ubiquitous. Kissing passionately entails kissing. Walking while chewing gum entails walking. Driving under 65 mph entails driving under 100 mph. And stretching a t_1 and then going for a run at t_2 entails going for a run at t_2 .

Now, suppose that I’ve promised to bring one of my famous home-made pumpkin pies to my family’s Thanksgiving celebration. It seems, then, that I have a reason to bake a pie, for doing so is a necessary means to my fulfilling this promise. And it seems that I have a reason to bake a pumpkin pie, for, again, doing so is a necessary means to my fulfilling this promise. This raises the question, though: What is the relationship between the reason that I have to bake a pie and the reason that I have to bake a pumpkin pie? More specifically,

¹ Options needn’t be mutually exclusive. I can have both the option of baking a pumpkin pie and the option of baking a pie. Only *alternative* options must be mutually exclusive. Thus, ϕ and ψ are alternative options if and only if ϕ -ing is an option, ψ -ing is an option, but both ϕ -ing and ψ -ing is not an option.

which, if either, is more fundamental than the other? Do I have a reason to bake a pie because I have a reason to perform some instance of pie-baking—e.g., pumpkin-pie baking? Or do I have a reason to bake a pumpkin pie because I have a reason to bake a pie? Or are they equally fundamental, as they would be if, say, I had a reason to do each because each would have optimal consequences?

The aim of this paper is to compare two possible answers to this question—*omnism* and *maximalism*—and to argue that the latter is preferable. Roughly speaking, maximalism is the view that only those options that are not entailed by any other option are to be assessed in terms of whether they have some feature (such as that of having optimal consequences), whereas omnism is the view that all options are to be assessed in terms of whether they have this feature. I argue that there are at least two reasons to prefer maximalism, for it is able to overcome two critical problems with omnism. But, first, I'll explain the two views more thoroughly.

1. Omnism and Maximalism about Reasons

On maximalism, we must distinguish between two types of options: maximal and non-maximal options. A maximal option is an option that either is not entailed by any other option or is entailed only by normatively equivalent options.² Any option that is not a maximal option is a non-maximal option. Admittedly, I've hedged a bit, here, in saying that a maximal option is either one that is not entailed by any other option or one that is entailed only by normatively equivalent options. I do so because it's not clear to me that an agent will always have at least one option that is not entailed by any other option. For instance, I may have all the following options: (Opt₁) thinking of a number greater than 1, (Opt₂) thinking of a number greater than 2, (Opt₃) thinking of a number greater than 3, and so on and so forth, *ad infinitum*. It may be, then, that for any option Opt_{*n*} there is always another option (*viz.*, Opt_{*n*+1}) that entails it. And, in that case, every option that consists in thinking of a number will be entailed by some other option. But I'm not entirely convinced. Even if, for every number *n*, there is always a greater number (e.g., *n*+1), it doesn't follow that I always have

² Two options are distinct options if and only if it is not the case that each entails the other. For instance, I have the option of typing as well as the *other* (that is, distinct) option of typing slowly. Typing slowly is a distinct option, because typing slowly is not the same as typing. After all, typing slowly is not entailed by typing. Of course, this is not to say that typing and typing slowly are *alternative* (as opposed to merely distinct) options. They're not. See note 1.

the option of thinking of that greater number. My abilities (and, hence, my options) are limited in a way that the set of numbers is not. So, perhaps, there are some numbers that I am unable to think of. I'm just not sure. But rather than try to resolve this difficult issue, it is sufficient for my purposes for us to think of a maximal option as an option that is entailed only by normatively equivalent options.³ To illustrate, suppose that I will win a million dollars if and only if I think of a number greater than 10. And assume that everything else is equal. In that case, thinking of 11 will be normatively equivalent to thinking of any of the following numbers: 12, 13, 14, and so on. And, so, my thinking of 11 will be a maximal option as will my thinking of 12, 13, or 14. Thus, in this case, I will have an infinite number of maximal options—assuming, that is, that there are no numbers that I unable to think of. This, then, is why I've hedged a bit in my definition of a maximal option. But even with this bit of hedging, the definition is perfectly coherent and so sufficient for my purposes.

We must distinguish between these two types of options, because, on maximalism, the reasons for performing a maximal option are fundamental and the reasons for performing a non-maximal option are derivative of the reasons for performing the options that entail it. More precisely, the view says the following.

Maximalism: (Max₁) For any subject S and any non-maximal option ν , S has a reason (or a reason of a certain type) to ν if and only if there exists an option φ such that S has a reason (or a reason of this type) to φ and S's φ -ing entails S's ν -ing. And, (Max₂) for any subject S and any maximal option μ , S has a reason (or a reason of a certain type) to μ if and only if S's μ -ing has feature F.⁴

³ The idea that we can think of a maximal option as one that entailed only by normatively equivalent options comes from Brown 2015.

⁴ Both here and in my formulation of omnism below, I leave implicit that the right-hand sides of these bi-conditionals have explanatory priority. Thus, we are to assume that, when the left-hand side of one of these bi-conditionals holds, it holds in virtue of the fact that the right-hand side holds.

Note also that my use of the term 'maximalism' is a bit idiosyncratic. Maximalism is typically taken to be an account of what one ought (or is permitted or is obligated) to do rather than an account of what one has reason to do. See, for instance, Brown 2015, Bykvist 2002, and Gustafsson 2014. But if, as many now think, what one ought (or is permitted or is obligated) to do is to be explained in terms of what one has reason to do, then we should think that this sort of maximalist view about what one ought to do must be combined with a maximalist view about what one has reason to do. And the above is an account of such a maximalist view, an account that I first introduced in Portmore 2013 but develop here.

On maximalism, if I have a reason to bake, this is because I have a reason to perform some instance of baking, such as baking a pie. And if I have a reason to bake a pie, this is because I have a reason to perform some instance of pie-baking, such as baking a pumpkin pie. And if I have a reason to bake a pumpkin pie, this is because I have a reason to perform some instance of pumpkin-pie-baking, such as baking a pumpkin pie and then taking it to my family's Thanksgiving celebration. And so on and so forth. But, of course, at some point we'll arrive at an option that either is not entailed by any other option or is entailed only by normatively equivalent options—that is, a maximal option.⁵ And when we do, the reason for performing this option cannot derive from the reason there is to perform some other option that entails it. For, given that it is a maximal option, there is no other option (or, at least, no normatively distinct option) that entails it. Thus, the reason for performing this option lies not with the fact that it is entailed by some other (normatively distinct) option, but with the fact that performing it has some reason-providing feature F. And 'F' can stand for whatever we take to be the most general reason-providing feature of options to be, including, perhaps, 'has optimal consequences', 'maximizes S's utility', or 'accords with Kant's categorical imperative'. Or it could be some long and complex disjunction, such as 'has one of the following normatively relevant properties: P_1, P_2, \dots, P_n '.

In contrast to maximalism, omnism holds that the reason for performing an option (be it a maximal or a non-maximal option) always lies with the fact that it has feature F. More precisely, the view is as follows.

Omnism: (Omni) For any subject S and any option ϕ , S has a reason (or a reason of a certain type) to ϕ if and only if S's ϕ -ing has feature F.

⁵ One may doubt that there will always be an option that is entailed only by normatively equivalent options. Suppose, for instance, that there is no limit to how large a number I can think of and that, for whatever number I do think of, God will give me and my loved ones precisely those many days in heaven. In that case, no matter what number n I think of, there will be an alternative option, Opt_{n+1} , that would be better in terms of the normatively relevant considerations. And so, in this case, there seems to be no option that is entailed only by normatively equivalent options. But this is a problem for the maximalism only if we think that I must always have sufficient reason to perform at least one of my options. After all, if there is no permissible maximal option that entails my opting for Opt_n , the maximalist will just claim that I don't have sufficient reason to opt for Opt_n . And, to my mind, that seems exactly right. For it seems that, in this case, I will do something that I lack sufficient reason to do no matter what I do, because no matter what I do there will be an option that I had more reason to perform, could have performed, but didn't perform. Thanks to Ángel Pinillos and Brad Armendt for pressing me on this.

Omnism and maximalism agree on reasons for performing maximal options: there is a reason to perform a maximal option if and only if it has feature F. But they disagree on reasons for performing non-maximal options. Whereas omnism holds that there is a reason to perform a non-maximal option if and only if it has feature F, maximalism holds that whether there is a reason to perform a non-maximal option depends, not on whether it has feature F, but on whether it is entailed by an option that has feature F. It's this difference that presents a problem for omnism, or so I'll now argue.

2. The Closure Problem for Omnism

Let's call a reason that is sufficient to ground the fact that S ought to φ when φ is no morally worse than any of S's other options in terms of any non-consequentialist considerations "a consequentially-decisive reason for S to φ ." Thus, such a reason would be sufficient to ground the fact that S ought to φ whenever either the only relevant considerations are consequentialist ones or the non-consequentialist reasons also favor S's φ -ing. And, now, consider the following three claims.

- (C1) *Having optimal consequences is a consequentially-decisive reason:* For any subject S and any option φ , S has a consequentially-decisive reason to φ if and only if S's φ -ing would have optimal consequences—that is, if and only if there is no alternative ψ such that the way things would be if S were to ψ is better than the way things would be if S were to φ .
- (C2) *Having a consequentially-decisive reason is closed under performance entailment:* For any subject S and any two options φ and ψ , if S has a consequentially-decisive reason to φ and S's φ -ing entails S's ψ -ing, then S has a consequentially-decisive reason to ψ .
- (C3) *Having optimal consequences is not closed under performance entailment:* There exists a subject S_1 with both the option of X-ing and the option of Y-ing such that: S_1 's X-ing would have optimal consequences, S_1 's X-ing entails S_1 's Y-ing,

but S_1 's Y-ing would not have optimal consequences.⁶

The omnist must accept C1. After all, the omnist holds that, for any subject S, any option φ , and any type of reason T, S has a T-type reason to φ if and only if S's φ -ing has feature F. And, of course, one type of reason is a consequentially-decisive one. Thus, the omnist must hold that, for any subject S and any option φ , S has a consequentially-decisive reason to φ if and only if S's φ -ing has feature F. Now, either 'F' stands for 'would have optimal consequences' or it doesn't. If it doesn't, then she'll have to deny that subjects ever have a consequentially-decisive reason to perform the option that would have optimal consequences, which is quite implausible. So the omnist should insist that 'F' stands for 'would have optimal consequences'. And that means that, on omnism, a subject will have a consequentially-decisive reason to perform an option if and only if her doing so would have optimal consequences. That is, the omnist must accept C1.

But this creates a problem for the omnist, because, if C1 is true, then at least one of C2 and C3 must be false. This is because C1–C3 together entail a logical contradiction. Here's the proof.

1. S_1 's X-ing would have optimal consequences. [From C3]
2. For any subject S and any option φ , S has a consequentially-decisive reason to φ if and only if S's φ -ing would have optimal consequences. [C1]
3. Therefore, S_1 has a consequentially-decisive reason to X. [From 1 and 2]
4. S_1 's X-ing entails S_1 's Y-ing. [From C3]
5. For any subject S and any two options φ and ψ , if S has a consequentially-decisive reason to φ and S's φ -ing entails S's ψ -ing, then S has a consequentially-decisive reason to ψ . [C2]
6. Therefore, S_1 has a consequentially-decisive reason to Y. [From 3, 4, and 5]
7. S_1 's Y-ing would not have optimal consequences. [From C3]
8. Therefore, it is not the case that S_1 has a consequentially-decisive reason to Y. [From 2 and 7]
9. Therefore, S_1 has a consequentially-decisive reason to Y, and it is not the case that

⁶ This is the reasons analogue of what's called the problem of act versions. See Brown 2015.

S_1 has a consequentially-decisive reason to Y . [From 6 and 8]

This is a problem for the omnist, because it's very difficult to deny either C2 or C3. Take C2, first. It seems that we need C2 to account for the validity of the following seemingly valid inference form: (1) S has a consequentially-decisive reason to φ . (2) S 's φ -ing entails S 's ψ -ing. Therefore, (3) S has a consequentially-decisive reason to ψ . To illustrate, suppose that I'm likely to get a ticket unless I drive under 65 mph and that my getting a ticket would have suboptimal (indeed, bad) consequences. In that case, I have a consequentially-decisive reason to drive under 65 mph. And, given that driving under 65 mph entails my driving under 100 mph, it seems that I have a consequentially-decisive reason to drive under 100 mph.⁷ So, in order to account for the validity of this and other similar inferences, we should, I believe, accept C2.

But what about C3? To see that it's true, consider the following example:

Two Drugs: Dr. Singh is deliberating at t_0 about what drugs, if any, to give her patient, Patrick, at t_2 . In fact, there are exactly two drugs available to her: A and B. If she gives Patrick both A and B at t_2 , using her right hand to inject him with A and her left hand to inject him with B, he'll be cured of his painful, but non-fatal, medical condition. If she gives him just one of A and B at t_2 , he'll die immediately. And if she gives him neither A nor B at t_2 , he'll be rendered incurable, having, then, to live the rest of his life with this painful medical condition. And although this would be quite bad, it would not be as bad as his dying immediately. Now, as a matter of fact, Dr. Singh is going to form at t_1 the intention to kill him by giving him just A at t_2 . Moreover, she's going to follow through with this intention. For, as it turns out, she has a grudge against Patrick. Given that she's going to form, and then follow through, with the intention to give him just A at t_2 , the following subjunctive conditional is true: if she were to give him A at t_2 , she would not give him B at t_2 . Nevertheless, assume that, if she were to respond appropriately to the decisive

⁷ Someone may object: (1) One way of driving under 100 mph is to drive exactly 80 mph. And (2) if I have a consequentially-decisive reason to drive under 100 mph, and driving exactly 80 mph is a way of driving under 100 mph, then I must have a consequentially-decisive reason to drive exactly 80 mph. But (3) I don't have a consequentially-decisive reason to drive exactly 80 mph. Therefore, (4) I don't have a consequentially-decisive reason to drive under 100 mph, as C2 implies. But we should reject (2). For an explanation of this as well as a counterexample to (2), see Nair 2015.

reason she has for curing him and thereby come instead to form at t_1 the intention to cure him by giving him both A and B at t_2 , she would follow through with this intention and Patrick would live happily ever after.

Clearly, Dr. Singh should give Patrick both A and B at t_2 , as this would have optimal consequences. It would, after all, cure Patrick of his painful medical condition, and curing him would have better consequences than that of any alternative—the alternatives being killing him (by giving him either only A or only B at t_2) and rendering him incurable (by giving him neither A nor B at t_2). But, interestingly, Dr. Singh's giving Patrick A at t_2 would not have optimal consequences. Indeed, it would have terrible consequence. For it's stipulated that, if she were to give him A at t_2 , she would not give him B at t_2 . Thus, her giving him A at t_2 would result in his death. Yet, her giving him both A and B at t_2 entails her giving him A at t_2 . So, *Two Drugs* establishes C3: There exists a particular subject S_1 (viz., Dr. Singh) with both the particular option of X-ing (i.e., giving Patrick both A and B at t_2) and the particular option of Y-ing (giving Patrick A at t_2) such that: S_1 's X-ing would have optimal consequences, S_1 's X-ing entails S_1 's Y-ing, but S_1 's Y-ing would not have optimal consequences.

So the omnist must either deny that *Two Drugs* is a coherent case or deny the validity of the sorts of inferences that C2 validates. Both are hard bullets to bite. It would be better, then, to deny C1, which we can, and must, do if we accept maximalism. For if we accept maximalism, we must accept the following instead of C1.

(C1*) *Having optimal consequences is a consequentially-decisive reason only for maximal options:* (a) For any subject S and any non-maximal option ν , S has a consequentially-decisive reason to ν if and only if there exists an option φ such that S has a consequentially-decisive reason to φ and S's φ -ing entails S's ν -ing. And, (b) for any subject S and any maximal option μ , S has a consequentially-decisive reason to μ if and only if S's μ -ing would have optimal consequences.

The problem with omnism's C1 is that it implies both that Dr. Singh has a

consequentially-decisive reason to give Patrick both A and B at t_2 and that she doesn't have a consequentially-decisive reason to give Patrick A at t_2 . This, in conjunction with C2 and the fact that the former entails the latter, generates a contradiction—see the argument above. But maximalism avoids implying a contradiction, because maximalism's C1* implies that Dr. Singh not only has a consequentially-decisive reason to give Patrick both A and B at t_2 but also to give him A at t_2 , and this despite the fact that her giving him A at t_2 would not itself have optimal consequences. For, on maximalism, whether she has a consequentially-decisive reason to give him A at t_2 depends, not on whether it would have optimal consequences, but on whether it is entailed by an option that would have optimal consequences. And it is entailed by an option that would have optimal consequences, because the option of giving him both A and B at t_2 would have optimal consequences and this option entails giving him A at t_2 .

So one reason to favor maximalism is that, in contrast to omnism, it is compatible with both C2 and C3. But this isn't the only reason to favor maximalism over omnism. There is, as I'll show in the next section, another instance where the omnist is committed to a claim that when combined with other plausible claims generates a contradiction.

2. The Moral Problem for Omnism

So far, I've specified omnism's and maximalism's views only on when there is *a reason* to perform an option. But that's not enough. For suppose that a subject S has both a reason to φ and a reason to ψ , but can't do both. If we want to account for what she ought to do in terms of her reasons, we'll need to account for whether she has *more reason* to φ , more reason to ψ , or neither more reason to φ than to ψ nor more reason to ψ than to φ . Thus, we'll need to supplement the above with the following.

Omnism (continued...): (Omn₂) For any subject S and any two options φ and ψ , S has more reason (or more reason of a certain type) to φ than to ψ if and only if she ought (an ought of that same type) to prefer how things would be if she were to φ to how things would be if she were to ψ .⁸

⁸ If there is not just one way that the world would be if S were to φ but only several different ways that the world could be if S were to φ , then we should replace all talk of the way the world be if S were to φ with talk of the prospect of S's φ -ing. The prospect of S's φ -ing is a probability distribution over the various possible ways that the world could turn out if S were to φ .

Maximalism (continued...): (Max₃) For any subject S and any two maximal options μ_x and μ_y , S has more reason (or more reason of a certain type) to μ_x than to μ_y if and only if she ought (an ought of that same type) to prefer how things would be if she were to μ_x to how things would be if she were to μ_y . (Max₄) For any subject S and any two non-maximal options ν_x and ν_y , S has more reason (or more reason of a certain type) to ν_x than to ν_y if and only if she ought (an ought of that same type) to prefer how things would be if she were to perform an optimal maximal option that entails her ν_x -ing to how things would be if she were to perform an optimal maximal option that entails her ν_y -ing.

Now, there are two bits of terminology in need of clarification. First, whether an agent ought to prefer how things would be if she were to φ to how things would be if she were to ψ just depends on what the given version of omnism or maximalism takes 'has feature F' to stand for. Thus, if 'has feature F' stands for 'produces optimal utility', then S ought to prefer how things would be if she were to φ to how things would be if she were to ψ if and only if her φ -ing would produce more utility than her ψ -ing would. And, if 'has feature F' stands for 'accords with Kant's categorical imperative', then S ought to prefer how things would be if she were to φ to how things would be if she were to ψ if and only if either only her ψ -ing would violate Kant's categorical imperative or they both would (or would not) but her ψ -ing would be a worse a violation (or a better non-violation). Thus, the above two formulations are meant to be substantively neutral such that an agent ought always to prefer her performing a permissible act to an impermissible act and ought always to prefer her performing a better permissible/impermissible act to a worse permissible/impermissible act.⁹ Second, a maximal option μ_x is an optimal maximal option if and only if there is no alternative maximal option μ_y such that the subject has more reason (or more reason of the type in question) to μ_y than to μ_x .

Having clarified these two bits of terminology, I can now present the moral problem for omnism—the problem being that the omnist cannot accept a plausible moral theory. To

That is, it is the set of mutually exclusive and jointly exhaustive ways the world might turn out if S were to φ , with each possibility assigned a probability such that the sum of the probabilities equals 1.

⁹ See Portmore 2011 for how this could be done for any sort of substantive view.

illustrate the problem, consider the following four claims.

- (C4) One has more moral reason to save ten distant strangers from death than to refrain from physically threatening someone.
- (C5) The self-interested reason that one has to buy a new car for oneself morally justifies acting contrary to the moral reason that one has to save ten distant strangers from death.
- (C6) For any subject S and any two options φ and ψ , if S has more moral reason to φ than to ψ , then whatever reason would morally justify her acting contrary to the moral reason that she has to φ would also morally justify her acting contrary to the moral reason that she has to ψ .
- (C7) The self-interested reason that one has to buy a new car for oneself does not morally justify acting contrary to the moral reason one has to refrain from physically threatening someone.

It seems that the omnist must accept C4 if she wants her moral theory to be plausible. To see why, consider the following case.

Roy: Roy knows that Terry has just inherited a substantial amount of money from a wealthy relative. Were Roy to intimidate Terry into giving him the money by physically threatening him, Roy would then donate the money so obtained to Oxfam International, thereby saving ten people from death. If, however, Roy doesn't physically threaten Terry, Terry is going to use the money to buy himself a new car, and those ten people that Roy could have saved will instead die. Everything else, though, will be the same. Thus, we are to assume that Roy would avoid punishment even if he were to threaten and steal from Terry.¹⁰

It seems that, on any plausible moral theory, there will be some number of people (n) such that one morally ought to prefer, other things being equal, how things would be if one were to physically threaten someone so as to obtain the money that one would then use to

¹⁰ This case is closely modeled after the case of Stan from Dorsey 2013, p. 366.

save n or more people from death to how things would be if one were to refrain from doing so and let those n or more people die. Moreover, it seems that, on any plausible moral theory, n will be less than or equal to ten. So even though there will certainly be plausible moral theories that include an agent-centered constraint against either physically threatening or stealing from someone, it seems that, to be plausible, such theories must hold that this constraint has a threshold such that, if enough is at stake, one will be permitted to perform the otherwise prohibited act-type. And, given Omniz, the omnist must accept C4 if she holds that one morally ought to prefer, other things being equal, how things would be if one were to physically threaten someone so as to save ten to how things would be if one were to refrain from doing so and let those ten die. And it seems that one morally ought to have this preference given that one should care more about saving the lives of ten people than keeping one's hands clean of theft and threats. Thus, it seems that the omnist must accept C4 if she wants her moral theory to be plausible. And, thus, she'll have to accept that Roy is permitted to physically threaten Terry so as to get the money to save the ten.

But this creates a problem for the omnist, because, if C4 is true, then at least one of C5–C7 must be false. This is because C4–C7 together entail a logical contradiction. Here's the proof.

1. One has more moral reason to save ten distant strangers from death than to refrain from physically threatening someone. [C4]
2. The self-interested reason that one has to buy a new car for oneself morally justifies acting contrary to the moral reason that one has to save ten distant strangers from death. [C5]
3. For any subject S and any two options φ and ψ , if S has more moral reason to φ than to ψ , then whatever reason would morally justify her acting contrary to the moral reason that she has to φ would also morally justify her acting contrary to the moral reason that she has to ψ . [C6]
4. Therefore, the self-interested reason that one has to buy a new car for oneself morally justifies acting contrary to the moral reason one has to refrain from physically threatening someone. [From 1–3]
5. It is not the case that the self-interested reason that one has to buy a new car for

oneself morally justify acting contrary to the moral reason one has to refrain from physically threatening someone. [From C7]

6. Therefore, the self-interested reason that one has to buy a new car for oneself morally justifies acting contrary to the moral reason one has to refrain from physically threatening someone, and it is not the case that the self-interested reason that one has to buy a new car for oneself morally justifies acting contrary to the moral reason one has to refrain from physically threatening someone. [From 4 and 5]

This is a problem for the omnist, because it's very difficult to deny any of C5–C7.¹¹ Take C5, first. To deny C5 is to accept that morality is extremely demanding—at least, assuming, as many philosophers believe, that one could easily save the lives of ten distant strangers by donating many thousands of dollars (that is, the cost of a new car) to certain charities.¹² Now, it's one thing to hold that we have an imperfect duty of beneficence that requires each of us “to make the happiness of others a serious, major, continually relevant, life-shaping end” (Hill 2002, p. 206) and, thus, to act throughout our adult lives so as to exhibit a sufficient propensity to promote this end—perhaps, by dedicating at least 20% of our time and resources to doing so.¹³ But it's quite another to hold that we are required to act in a way that sacrifices our time and resources for the sake of saving the lives of distant strangers each and every time we morally ought to prefer how things would be if we did to how things would be if we didn't. For such a theory would require us to dedicate much more than just 20% of our time and resources to promoting the happiness of others. And that would be extremely demanding, leaving us with little to no time or resources to further our own happiness. And I don't believe that a plausible moral theory could be so demanding.¹⁴ To illustrate, consider the following case.

“Gus: Gus finds himself the recipient of an inheritance from a wealthy relative. This inheritance will allow Gus to buy a new car, which Gus desires to do, and which will allow Gus to see much more of his significant other, who lives in a distant town.

¹¹ This problem draws inspiration from discussions in Archer 2015, Dorsey 2013, Kagan 1989, and Kamm 1985.

¹² See, for instance, Singer 2009 and Unger 1996.

¹³ For more on this, see Noggle 2009.

¹⁴ For a further defense of this claim, see Portmore 2011, especially chapter two.

Alternatively, Gus could donate his inheritance to Oxfam International, which would save ten people from death.” (Dorsey 2013, p. 365)

A plausible moral theory would not, I believe, imply that Gus is obligated to donate his inheritance to Oxfam International. A plausible moral theory may hold that Gus is obligated to dedicate substantial portions of his time and resources to helping others, but it would not hold that he is required to do so on every occasion in which the outcome of his doing so is to be preferred, morally speaking, to that of his not doing so. So unless the omnist wants to accept an implausibly demanding moral theory, she’ll need to endorse C5.

C6 is, I believe, even more difficult to deny, for it seems analytically true. If some reason R is able to morally justify acting contrary to the moral reason that one has to φ , which is weightier than the moral reason one has to ψ , then mustn’t it also be able to morally justify acting contrary to the weaker moral reason that one has to ψ ? Thus, if the omnist wants to deny C6, she owes some explanation of why we should think that this seemingly analytical truth is false.

This leaves us with C7, which is also very difficult to deny. To deny C7 is to accept a moral theory that’s rather egoistic and, thus, overly permissive. Such a moral theory would permit one to harm someone and violate his or her autonomy merely for the sake of personal gain. To illustrate, consider the following case.

“Stan: Stan knows that Jerry has just inherited a substantial amount of money from a wealthy relative. Were Stan to intimidate Jerry into giving him the money as a result of [physically threatening] Jerry..., this would be a prudential benefit to Stan, given that this would allow him to buy a new car, from which he will derive pleasure, and which will allow Stan to see much more of his significant other, who lives in a distant town.” Everything else, though, will be the same. Thus, we are to assume that Stan would avoid punishment even if he were to threaten and steal from Jerry. (Dorsey 2013, p. 366)

But, on any plausible moral theory, Stan would not be permitted to threaten and steal from Jerry just so that he could use the money thereby obtained to buy a new car for

himself. Thus, the omnist has a problem—a moral problem. Since C6 is analytically true, she must deny C4, C5, or C7. And regardless of which she rejects, she'll be forced to accept an implausible moral theory. If she rejects C4, she must implausibly deny that Roy morally ought to prefer how things would be if he were to threaten Terry so as to save the ten distant strangers to how things would be if he were to refrain from threatening Terry and let those ten die. If she rejects C5, she'll have to hold that morality is overly demanding, requiring Gus to donate his inheritance to Oxfam International. And, if she rejects C7, she must hold that morality is overly permissive, allowing Stan to threaten and steal from Jerry just so he can buy a new car for himself. So if we want to accept a plausible theory of morality, we must deny omnism.

Fortunately, if we accept maximalism, we can (and, indeed, must) deny C4 and, thus, should deny that Roy has more moral reason to save the ten distant strangers from death than to refrain from physically threatening Terry. On maximalism, we don't look at whether Roy morally ought to prefer how things would be if he were to threaten Terry so as to save the ten to how things would be if he were to refrain from threatening him and let the ten die. Instead, we look at whether Roy morally ought to prefer how things would be if he were to perform an optimal maximal option that entails his threatening Terry to save the ten to how things would be if he were to perform an optimal maximal option that entails his refraining from threatening him and letting the ten die. And, as I'll now argue, Roy morally ought not to prefer the former to the latter. For it seems that any optimal maximal option must strike a balance between using more of one's limited time and resources to promote the happiness of others and using more of one's limited time and resources to promote various other moral goals, such as justice, self-improvement, and the improvement of one's family. Thus, it seems that the optimal maximal options available to Roy will not be the ones in which Roy dedicates all of his time and resources to promoting the happiness of others, doing so at the expense of neglecting other important moral ends, such as justice, self-improvement, and improvement of his family. So, Roy will not have most moral reason to save as many distant strangers as he could save if he were to set aside these other important moral goals. For we could all save more distant strangers if we were to disregard justice, our families, and our own self-improvement. We could all save many more lives by, say, robbing the rich and donating the proceeds of these robberies to Oxfam International. But that's not what we have

most moral reason to do. It seems, then, that Roy will have most moral reason to save the lives of only n distant strangers over the course of his life, where n is significantly less than the number of distant strangers that he could save if he were to disregard all other moral goals.

Admittedly, Roy could save more than n strangers if he were to neglect these other moral goals. But it would not be morally better for him to do so, and so he should not prefer that he does so. After all, n is, by definition, the number that strikes the best balance between the goal of saving more strangers and the goal of doing more to promote various other moral ends. Moreover, if Roy were to insist on saving $n+10$ strangers, the best way for him to do so would not be by setting aside the moral goal of respecting justice (by, say, threatening Terry), but instead by setting aside some prudential goal (by, say, buying a used car instead of a new one) or by setting aside some less important moral goal (by, say, spending less time improving himself). It seems, then, that if Roy were to insist on saving $n+10$ strangers, it would be better for him to do so by using less of his time and resources on either his self-improvement or his own happiness than by obtaining additional resources through threatening and stealing from Terry.¹⁵ And, if all this is right, we should not think that Roy morally ought to prefer how things would be if he were to perform an optimal maximal option that entails his threatening Terry (one in which he saves the lives of $n+10$ strangers over the course of his life and threatens Terry so as to save those additional ten) to how things would be if he were to perform an optimal maximal option that entails his refraining from threatening Terry (one in which he saves the lives of only n strangers over the course of his life but neither threatens nor steals from anyone). And, so, we should think that Roy does not have more moral reason to threaten Terry and save the additional ten distant strangers than to refrain from threatening Terry and save only n distant strangers over the course of his life. Thus, we should deny that one has more moral reason to save ten distant

¹⁵ In the case of *Roy*, we are to assume that everything else will be the same regardless of whether or not he threatens Terry. Thus, we are to assume that Roy will, in the future, dedicate the same amount of time and resources to each moral and prudential goal that he has regardless of whether he threatens Terry and saves the ten at present. And this is why the omnist holds that Roy has more moral reason to threaten Terry than to refrain from doing so, for how things would be if he were to threaten Terry is to be preferred, morally speaking, to how things would be if he were not to do so. But, on maximalism, we don't look at how things would be if he were, at present, to act this way or that way. Instead, we look at how things would be if he were to perform this or that maximal option over the course of his future. And we derive his reasons for performing this or that action at present in terms of what those maximal options entail his doing at present.

strangers from death than to refrain from physically threatening someone. And, so, we should deny C4—or, at least, we should so long as we adopt Max₄ as opposed to Omni.¹⁶

3. An Objection to Maximalism

At this point, readers may wonder whether they must accept maximalism even if I'm right about these problems with omnism. That is, they may wonder whether omnism and maximalism exhaust the possibilities. In fact, they don't. We could, for instance, distinguish between minimal options (i.e., options that do not entail any other option) and non-minimal options (options that do entail some other option) and adopt the following view.

Minimalism: (Min₁) For any subject S and any non-minimal option ν , S has a reason (or a reason of a certain type) to ν if and only if there exists an option φ such that S has a reason (or a reason of this type) to φ and S's ν -ing entails S's φ -ing. And, (Min₂) for any subject S and any minimal option μ , S has a reason (or a reason of a certain type) to μ if and only if S's μ -ing has feature F.

But this view is a non-starter. On this view, I have a reason to feed my daughter shards of glass given that I have a reason to feed her and my feeding her shards of glass entails feeding her. That's absurd. Nevertheless, there could be some other view that I have yet to consider that is preferable to all three views. Whether there is some more plausible alternative view is not something that I have space here to consider. Yet I needn't do so. For my thesis is merely that maximalism is preferable to omnism, not that maximalism is true. And this is itself a significant thesis given that omnism is the current philosophical orthodoxy. The theories of reasons that philosophers have offered to date are all versions of omnism. Consider the following three illustrations. On subjectivism, S has a reason to φ if and only if S's φ -ing would advance some element in S's "subjective motivational set" (Williams 1981). On expected utility theory, S has decisive reason to φ if and only if S's φ -ing would maximize S's expected utility. And, on act-utilitarianism, S has sufficient moral reason to φ if and only if S's φ -ing would maximize aggregate utility. These are all versions

¹⁶ For more on this, see Portmore 2015c.

of omnism as they do not even distinguish between different sorts of options let alone give different criteria for assessing different sorts of options.

Of course, I can't establish even my modest thesis that maximalism is preferable to omnism merely by showing, as I believe I have, that omnism is subject to two critical problems to which maximalism is immune. For it could be that maximalism is subject to equally serious problems to which omnism is immune. And, in that case, maximalism wouldn't be preferable to omnism. So let's see if maximalism is subject to any equally serious problems.

One worry is that maximalism seems to be subject to clear counterexamples. To illustrate, suppose that I have a reason to relieve some of my current stress and that I can do so by presently engaging in any manual activity—that is, any activity that involves working with my hands. And assume that, with regard to relieving my stress, it doesn't matter what the activity is so long as it's one that involves working with my hands. In that case, I have a reason, R_{WH} , to *work with my hands*, because doing so would relieve some of my stress. What's more, I have a reason, R_{SS} , to *sculpt a statue* out of clay, because doing so would relieve some of my stress. And my sculpting a statue out of clay (that is, SS) entails my working with my hands (that is, WH). Yet, it doesn't seem as though R_{WH} derives from R_{SS} . Rather, it seems that I have R_{SS} because I have R_{WH} —that is, it seems that I have a reason to sculpt a statue out of clay because I have a reason to work with my hands. So isn't this a clear counterexample to maximalism?¹⁷

No, it is not. For maximalism doesn't imply that R_{WH} derives from R_{SS} . It implies only that R_{WH} derives from a reason to perform some option that entails working with my hands. And this seems absolutely right. I have a reason to work with my hands because I have a reason to perform a maximal option that entails my working with my hands. After all, if there were no maximal option that I both had a reason to perform and that entailed my working with my hands, then I don't see how I could have a reason to work with my hands. Moreover, maximalism allows that I have a reason to sculpt a statue out of clay because I have a reason to do some activity (indeed, any activity) that involves working with my hands. I have a reason to sculpt a statue out of clay because I have a reason to perform a maximal option that involves my presently engaging in some manual activity. Thus, the fact

¹⁷ I thank an anonymous reviewer from pressing this worry.

that I can relieve some of my current stress by presently engaging in any manual activity is a reason for me to perform a maximal option that involves my presently engaging in any manual activity. And the fact both that I have a reason to perform a maximal option that involves my presently engaging in any manual activity and that, by presently sculpting a statue out of clay, I can ensure that I perform a maximal option that involves my presently engaging in some manual activity is a reason for me to sculpt a statue out of clay. So I don't see that there are any problems with maximalism. In any case, I challenge others to show that there are problems with maximalism that are just as great as the two problems that I've identified for omnism.

4. Conclusion

As we've seen, omnism is subject to two significant problems. First, the omnist's account of when one has a reason to φ (i.e., Omni_1) commits her to C1, which is problematic because it's incompatible with her accepting both C2 and C3. And because of this the omnist must either deny the coherence of *Two Drugs* or the validity of the inferences that C2 validates. Both are hard bullets to bite. Second, the omnist's account of when one has more moral reason to φ than to ψ (i.e., Omni_2) commits her to C4, which is problematic because it's incompatible with her accepting all of C5–C7. And because of this and the fact that C6 seems analytic, the omnist must accept an implausible moral theory, one that implausibly holds that morality is overly demanding, overly permissive, or committed to the view that Roy morally ought not to prefer how things would be if he were to threaten Terry so as to save the ten to how things would be if he were to refrain from doing so and let the ten die. But if we accept maximalism we avoid both problems. We have, therefore, at least two good reasons for preferring maximalism to omnism. And so absent any reason to think that maximalism suffers from problems at least as great as these two that I've identified with omnism (and I see none), we should prefer maximalism to omnism. And if that's right, then most of the theories of reasons that have been offered to date are unsatisfactory as they stand.

Arizona State University
dwportmore@gmail.com

References

- Archer, A. (2015). "Moral Obligation, Self-Interest and The Transitivity Problem." Working manuscript.
- Brown, C. (2015). "Consequentialism, Maximalism, and the Structure of Acts." Working manuscript.
- Bykvist, K. (2002). "Alternative Actions and the Spirit of Consequentialism." *Philosophical Studies* 107: 45–68.
- Dorsey, D. (2013). "The Supererogatory, and How to Accommodate It." *Utilitas* 25: 355–382.
- Gustafsson, J. E. (2014). "Combinative Consequentialism and the Problem of Act Versions." *Philosophical Studies* 167: 585–596.
- Hill, Jr., T. E. (2002). *Human Welfare and Moral Worth*. Oxford: Oxford University Press.
- Kagan, S. (1989). *The Limits of Morality*. Oxford: Oxford University Press.
- Kamm, F. M. (1985). "Supererogation and Obligation." *The Journal of Philosophy* 82: 118–138.
- Nair, S. (2015). "Conflicting reasons, unconflicting 'ought's." *Philosophical Studies*. DOI 10.1007/s11098-015-0511-4.
- Noggle, R. (2009). "Give 'Till It Hurts? Beneficence, Imperfect Duties, and a Moderate Response to the Aid Question." *Journal of Social Philosophy* 40: 1–16.
- Portmore, D. W. (2015a). "Maximal Moral Harmony." Working manuscript that's available at <http://bit.ly/1NmZ57E>.
- . (2015b). "Morality, Rationality, and Performance Entailment." Working manuscript that's available at <http://bit.ly/1L6iVOI>.
- . (2015c). "Transitivity, Moral Latitude, and Supererogation." Working manuscript that's available at <http://bit.ly/1laIFTP>.
- . (2013). "Perform Your Best Option." *The Journal of Philosophy* 110: 436–459.
- . (2011). *Commonsense Consequentialism: Wherein morality meets rationality*. New York: Oxford University Press.
- Singer, P. (2009). *The Life You Can Save*. New York: Random House.
- Unger, P. (1996). *Living High and Letting Die*. Oxford: Oxford University Press.
- Williams, B. (1981). "Internal and External Reasons." In his *Moral Luck*, pp. 101–113. Cambridge: Cambridge University Press.