

Maximalism vs. Omnism about Permissibility*

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 both of these options are permissible. This raises the issue of which, if either, is more fundamental than the other. Is baking a pie permissible because it's permissible to perform some instance of pie-baking, such as pumpkin-pie baking? Or is baking a pumpkin pie permissible because it's permissible to bake a pie? Or are they equally fundamental, as they would be if they were both permissible because, say, they both have optimal consequences? The aim of this paper is to compare two alternative responses to this issue—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 right-making feature *F* (such as that of having optimal consequences), whereas omnism is the view that all options are to be assessed in terms of whether they are *F*. I argue that maximalism is preferable to omnism because it provides a more plausible solution to the problem of act versions and is not subject to any problems of its own. And if I'm right about maximalism's being preferable to omnism, then most moral theories, which are all versions of omnism, need significant revision.

Keywords: closure principles, deontic logic, Gustafsson, Jackson, maximalism, morality, omnism, options, performance entailment, permissibility, problem of act versions, Professor Procrastinate.

The performance of one option can entail the performance of another. That is, performing the one can necessitate performing 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 I can't bake a pumpkin pie without baking.¹ More precisely, for any two of my options φ and ψ , my φ -ing entails my ψ -ing if and only if I don't have the option of φ -ing without ψ -ing. Such instances of performance entailment are ubiquitous. Kissing passionately entails kissing. Driving under 55 mph entails driving under 100 mph. Marrying an unmarried man entails marrying a bachelor. Typing the word 'the' entails typing the letter 't'. Charging my iPhone entails using a charging cable. Stretching a t_1 and then going for a run at t_2 entails going for a run at t_2 .

¹ Note, then, that options needn't be mutually exclusive. I can have both the option of baking and the option of baking a pumpkin pie even though these two are not mutually exclusive. Only *alternative* options need to be mutually exclusive. Thus, two options, φ and ψ , are alternative options if and only if both φ -ing and ψ -ing is not an option. By contrast, two options, φ and ψ , are *distinct* options if and only if it is not the case that each entails the other. So, since baking doesn't entail baking a pumpkin pie, these two are distinct options, though they are not alternative options.

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'm permitted to bake a pie, for I'm permitted (indeed, obligated) to fulfill this promise and doing so is a necessary means to my fulfilling it. And it seems that I'm permitted to bake a pumpkin pie, for, again, I'm permitted to fulfill my promise and doing so is a necessary means to my fulfilling it. This raises the question, though: What is the relationship between the permissibility of my baking a pie and the the permissibility of my baking a pumpkin pie? More specifically, which, if either, is more fundamental than the other? Is baking a pie permissible because it's permissible to perform some instance of pie-baking, such as pumpkin-pie baking? Or is baking a pumpkin pie permissible because it's permissible to bake a pie? Or are they equally fundamental, as they would be if they were both permissible because, say, they both 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 that the permissibility of an option that entails others is more fundamental than the permissibility of the options that it entails. Thus, on maximalism, only those options that are not entailed by any other option (i.e., maximal options) are to be assessed in terms of whether they have some right-making feature F, such as that of having optimal consequences. The rest are to be assessed, not in terms of whether they are F, but in terms of whether they are entailed by a maximal option that is F. By contrast, omnism is the view that the permissibility of an option that entails others and the permissibility of the options that it entails are equally fundamental. Thus, on omnism, all options are to be assessed in terms of whether they are F. I argue that maximalism is preferable to omnism because it provides a more plausible solution to the problem of act versions and is not subject to any problems of its own. And if I'm right about maximalism's being preferable to omnism, then most moral theories, which are all versions of omnism, need significant revision.

In section 1, I explain both omnism and maximalism more thoroughly. In section 2, I explain the problem of act versions. In section 3, I argue that maximalism offers a more plausible solution to this problem. And, in section 4, I argue that maximalism isn't subject to any problems of its own.

1. Omnism and Maximalism about Permissibility

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 (that is, distinct) option or is entailed only by normatively equivalent options. 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—that is, options that are equivalent in terms of all normatively relevant factors. 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 my thinking of a number will be entailed by some other option. Nonetheless, 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 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 either is not entailed by any other option or is entailed only by normatively equivalent options.² 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, 14, and so on. In this case, then, I will have an infinite number of maximal options—assuming, that is, that there are no numbers that I lack the option of thinking 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, which is sufficient for my purposes.

² The idea that we can think of a maximal option as one that is entailed only by normatively equivalent options comes from Brown 2015.

We must distinguish between these two types of options, because, on maximalism, the permissibility of a maximal option is more fundamental than the permissibility of the non-maximal options that it entails. More precisely, the view says the following.

Maximalism: (Max₁) For any subject S and any maximal option μ , S 's μ -ing is permissible if and only if S 's μ -ing has feature F , and when S 's μ -ing is permissible, this is in virtue of the fact that S 's μ -ing has feature F . And, (Max₂) for any subject S and any non-maximal option ν , S 's ν -ing is permissible if and only if there is a distinct option φ such that S 's φ -ing is permissible and S 's φ -ing entails S 's ν -ing, and when S 's ν -ing is permissible, this is in virtue of the fact that there is a distinct option φ such that S 's φ -ing is permissible and S 's φ -ing entails S 's ν -ing.³

On maximalism, if I'm permitted to bake, this is because I'm permitted to perform some instance of baking, such as baking a pie. And if I'm permitted to bake a pie, this is because I'm permitted to perform some instance of pie-baking, such as baking a pumpkin pie. And if I'm permitted to bake a pumpkin pie, this is because I'm permitted 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.⁴ And when we do, the permissibility of this option cannot derive from the permissibility of 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 permissibility of this maximal option lies not with the fact that it is entailed by some other (normatively distinct) option—it isn't,

³ Cf. Brown 2015, Bykvist 2002, and Gustafsson 2014. Maximalism was first formulated, though not under that name, in Goldman 1978.

⁴ It's possible that there won't always be a maximal option even in this sense. That is, it's possible that there won't be any 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 I had more reason to perform. 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 maximalism only if we think that I must always have a permissible option. Yet this seems like a situation where I have no permissible option. For no matter what I do, there will be an alternative option that I had more reason to perform, could have just as easily performed, but didn't perform. It seems to me, then, that maximalism gets the right answer in this sort of case: every option is impermissible, because, for each option, there is no maximal option (and, thus, no permissible maximal option) that entails it. Thanks to Ángel Pinillos and Brad Armendt for pressing me on this.

but with the fact that it has some right-making feature F. And ‘has feature F’ can stand for whatever we take to be the most fundamental right-making feature of options to be, including, perhaps, ‘has optimal consequences’, ‘maximizes S’s utility’, or ‘accords with Kant’s categorical imperative’. Or it could stand for some long and complex disjunction, such as ‘has at least one of the following properties: P_1, P_2, \dots, P_n ’.

In contrast to maximalism, omnism holds that the permissibility of an option (be it a maximal or a non-maximal option) always depends on whether it has right-making feature F, and never on whether it is entailed by a maximal option with right-making feature F. More precisely, the view is as follows.

Omnism: (Omni_i) For any subject S and any option φ , S’s φ -ing is permissible if and only if S’s φ -ing has feature F, and when S’s φ -ing is permissible, this is in virtue of the fact that S’s φ -ing has feature F.

Omnism and maximalism agree on what accounts for a permissible maximal option’s being permissible: its having feature F. But they disagree on what accounts for a permissible non-maximal option’s being permissible. Whereas omnism holds that a non-maximal option is permissible if and only if it has feature F, maximalism holds that a non-maximal option is permissible if and only if it is entailed by a maximal option that has feature F and, so, regardless of whether it is itself F. It’s this difference that presents a problem for omnism, or so I’ll argue. The problem is that omnism cannot offer as plausible a solution to what is known as the problem of act versions.

2. The Problem of Act Versions

The problem of act versions was first discovered as a problem for act-consequentialism.⁵ The problem lies with the fact that some options are versions (that is, instances) of other options, and sometimes these versions have very different consequences than the options that they are versions of. For instance, opening a bag of cookies and eating just one is a version (that is, an instance) of opening that bag of cookies—in other words, opening a bag of cookies and eating just one entails opening that bag of cookies. But even if opening that bag of cookies and eating

⁵ See Bergström (1966) and Castañeda (1968).

just one has optimal consequences, opening it may not. For it may be that, as a matter of fact, you would eat all the cookies and get sick if you were to open the bag. In such a case, act-consequentialism implies both that you are obligated to refrain from opening the bag of cookies (since opening the bag would have suboptimal consequences) and that you are obligated to open the bag of cookies and eat just one (since opening the bag and eating just one would have optimal consequences). And this seems problematic given that there's no way to fulfill both obligations.⁶ For there is no way to open the bag of cookies and eat just one while refraining from opening the bag.

Although the problem of act versions is typically seen to be a problem only for act-consequentialism, it is, I'll argue, a problem for any moral theory that holds that the goodness of an act's consequences is relevant to assessing its permissibility. (And what plausible moral theory doesn't hold this?) To illustrate the generality of the problem, I'll need to appeal to cases in which each of the agent's options with optimal consequences are also optimal in terms of *non-consequentialist factors*: factors that are possibly morally relevant but don't have to do with the goodness of the act's consequences. Non-consequentialist factors include justice, fidelity, self-improvement, respect for autonomy, and respect for individuals as ends. And I'll call cases in which the agent's options with optimal consequences are also optimal in terms of non-consequentialist factors *consequential cases*.

Consequential cases are a proper subset of all cases. In all cases, some of the agent's options have optimal consequences. But, in many cases, these options with optimal consequences are suboptimal in terms of some non-consequentialist factor. For instance, it may be that the only way for an agent to produce optimal consequences is to murder one innocent person, thereby preventing five others from each committing a comparable murder. This would not count as a consequential case, because murdering the one would count as treating her as a mere means to saving the greater number, and treating someone as a mere means is a possibly morally relevant factor that doesn't have to do with the goodness of its

⁶ For more on this, see Kiesewetter 2015 (pp. 929–934). What he has to say is along very similar lines to what I said in Portmore 2011 (pp. 181–183). The main difference is that whereas Kiesewetter's argument is based on the procedural rational requirement that S be such that, if she believes that she's obligated to φ , she intends to φ , my argument was based on the substantive rational requirement that S intend to φ if she both is obligated to φ and has the capacity to intend to φ . But both arguments argue against the plausibility of conflicting obligations on the basis of this sort of connection between being (or believing that one is) obligated to φ and having a substantive (or procedural) rational requirement to intend to φ along with the idea that rationality requires that S be such that, if she believes that she will not both φ and ψ , then she does not both intend to φ and intend to ψ .

consequences. Nevertheless, some cases are consequential cases. For instance, suppose that the only way for a doctor to produce optimal consequences is for her to perform the life-saving surgery that the patient has autonomously requested and that she is contractually obligated to perform. And assume that there is nothing at all morally suspect about the doctor's performing this surgery. This, then, would be a consequential case.

Now, in consequential cases, it seems that the only permissible options are those that would produce optimal consequences.⁷ For instance, it seems that the doctor's only permissible option is to perform the surgery. Thus, it seems that we should accept the following minimal sort of consequentialism.

- (C1) *Minimal Consequentialism (the view that consequences matter)*: For any consequential case with a subject S who has the option of ϕ -ing, S's ϕ -ing is morally permissible if and only if S's ϕ -ing would have optimal consequences—that is, if and only if there is no option ψ such that the amount of goodness that there would be if S were to ψ is greater than the amount of goodness that there would be if S were to ϕ .

But when we combine this with the following two plausible claims, we get a version of the problem of act versions.

- (C2) *Permissibility Closure (the view that permissibility is closed under performance entailment)*: For any case with a subject S who has both the option of ϕ -ing and the option of ψ -ing, if S's ϕ -ing is morally permissible and S's ϕ -ing entails S's ψ -ing, then S's ψ -ing is morally permissible.
- (C3) *Optimality Non-Closure (the view that there exists a consequential case where the optimality of consequences is not closed under performance entailment)*: There exists a consequential case with a subject S₁ who has both the option of X-ing and the option of Y-ing such that: S₁'s X-ing would have optimal consequences,

⁷ If you're skeptical about maximizing, just replace 'optimal consequences' with 'good enough consequences' throughout. The problem of act versions will still arise.

S_1 's X-ing entails S_1 's Y-ing, but S_1 's Y-ing would not have optimal consequences.

The problem is that C1–C3 together entail a logical contradiction. Here's the proof.

1. S_1 's X-ing would have optimal consequences. [From C3]
2. S_1 's X-ing is morally permissible if and only if S_1 's X-ing would have optimal consequences. And, likewise, S_1 's Y-ing is morally permissible if and only if S_1 's Y-ing would have optimal consequences. [From C1 and C3]
3. Therefore, S_1 's X-ing is morally permissible. [From 1 and 2]
4. S_1 's X-ing entails S_1 's Y-ing. [From C3]
5. If S_1 's X-ing is morally permissible and S_1 's X-ing entails S_1 's Y-ing, then S_1 's Y-ing is morally permissible. [From C2 and C3]
6. Therefore, S_1 's Y-ing is morally permissible. [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 's Y-ing is morally permissible. [From 2 and 7]
9. Therefore, S_1 's Y-ing is morally permissible, and yet it is not the case that S_1 's Y-ing is morally permissible, which is a logical contradiction. [From 6 and 8]

The problem seems to lie with C1, 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 is morally permitted to φ . (2) S's φ -ing entails S's ψ -ing. Therefore, (3) S is morally permitted to ψ . To illustrate, suppose that I'm permitted to kiss my partner passionately. And, given that my kissing her passionately entails my kissing her, it seems that we can infer that I'm permitted to kiss her.⁸ Moreover, the validity of this type of inference explains the validity of means-end reasoning. People reason, seemingly validly: (1) I ought (and, so, am permitted) to do X. (2) My Y-ing is a necessary means to my X-ing. Therefore, (3) I ought (and, so, am permitted) to do Y. This is just an instance of the type of reasoning that C1 validates, for, if my Y-ing is a necessary means to my X-ing, then I cannot X without Y-ing, which is just to say that my X-ing entails my Y-ing. So, in order to give a

⁸ We can't infer that I'm permitted to kiss her non-passionately. But we can infer that I'm permitted to kiss her.

unified account of the validity of such inferences, we should, I believe, accept C2.⁹

And C3 is impossible to deny—or, at least, it is if the following sort of case is coherent.

The Injection: Doc's patient, Pat, needs an injection of drug D, and the sooner he gets it, the better. Doc has the options of injecting him with a proper dose of D, injecting him with an improper dose of D, and not injecting him with D. The worst thing would be for Doc to inject him with an improper dose of D. That would kill Pat and result in Doc's going to prison. It would be better not to inject him with D at all; in which case, Doc would lose her medical license but stay out of prison, while Pat would suffer for a few days before going to another doctor who would then inject him with a proper dose of D. It would be much better, though, if Doc were just to inject Pat with a proper dose of D. Indeed, this would be best. For, in that case, everyone would live happily ever after. Now, as a matter of fact, if Doc were to inject Pat with D, she would inject Pat with an improper dose of D. This is not because of coercion or outside interference or anything like that, but because Doc is now forming the intention to kill Pat, who belongs to a race that she deeply despises.

Assume that I haven't left out any morally important details from the above description. Given that, it's clear that this a consequential case. Of her three options, Doc's only option with optimal consequences is injecting Pat with a proper dose of D.¹⁰ And this option is also optimal in terms of every other possibly morally relevant factor—that is, it's optimal in terms of non-consequentialist factors. It's optimal in terms of fidelity given that doing anything else would entail her violating the Hippocratic Oath and her contractual obligation to provide her patient with the standard of care. It's optimal in terms of respect for autonomy given that Pat has, we'll assume, autonomously requested to receive the

⁹ Some philosophers reject the validity of the above inference form—see, most notably, Jackson and Pargetter 1986. Nevertheless, even Jackson now accepts the validity of some inferences of this type. Specifically, he claims that, from the fact that he ought to raise both his left arm and his right arm at *t*, we can infer that he ought to raise his left arm at *t*—see Jackson 2014, pp. 645–646. But I see no good reason for thinking that the above type of inference is valid only when X and Y are cotemporaneous. Now, Jackson thinks that the case of Professor Procrastinate provides such a reason, but, as I show below, it does not.

¹⁰ It also seems to be Doc's only option with good enough consequences, which is why I said, in note 7, that the problem of act versions arises even if we substitute 'good enough consequences' for 'optimal consequences' throughout C1–C3.

standard of care. And injecting him with a proper dose of D seems to be unsurpassed in terms of every other possibly relevant consideration. So it's a consequential case.

Given that it's a consequential case, minimal consequentialism implies that Doc is obligated to inject Pat with a proper dose of D, for this would produce better consequences than any alternative would.¹¹ But, unfortunately, minimal consequentialism also implies that Doc is obligated to refrain from injecting him with D, because doing so is suboptimal in terms of its consequences. For the consequences would be better if she were instead to refrain from injecting him with D. It's stipulated, after all, that if Doc were to inject Pat with D, she would inject him with an improper dose of D, which would result in Pat's dying and Doc's going to prison. And this would be much worse than what would happen if she were to refrain from injecting him with D: Doc's losing her medical license but staying out of prison and Pat's suffering longer but continuing to live. Thus, minimal consequentialism (that is, C1) implies both that Doc is obligated (and so permitted) to inject Pat with a proper dose of D and that Doc is obligated to refrain from injecting him with D (and so not permitted to inject him with D), and yet Doc cannot inject him with a proper dose of D without injecting him with D. Injecting him with a proper dose of D entails injecting him with D.

Now, these implications of minimal consequentialism in *The Injection* conflicts with C2, which implies that if Doc is permitted to inject Pat with a proper dose of D, then she must also be permitted to inject him with D given that her injecting him with a proper dose of D entails her injecting him with D. Thus, it's not just act-consequentialism, but also minimal consequentialism (or C1), that is subject to the problem of act versions. And this is significant, because it seems that most moral theories endorse minimal consequentialism. Yet, given the truth of C2 and C3, the only way to solve the problem of act versions is to deny C1 and its implication that, in *The Injection*, Doc is permitted to inject Pat with a proper dose of D but not permitted to inject him with D. Technically, we can do this either by denying that Doc is permitted to inject Pat with a proper dose of D or by denying that she's

¹¹ Act-consequentialism holds that, for any subject S and any option φ , S's φ -ing is permissible if and only if there is no alternative option ψ such that the amount of goodness that there would be if S were to ψ is greater than the amount of goodness that there would be if S were to φ . This in conjunction with the claim that S's φ -ing is obligatory if and only if φ -ing is S's only permissible option entails that, on act-consequentialism, S's φ -ing is obligatory if and only if, for every alternative option ψ , the amount of goodness that there would be if S were to φ is greater than the amount of goodness that there would be if S were to ψ .

not permitted to inject him with D. But it seems absurd to deny that Doc is permitted to inject Pat with a proper dose of D, for this is clearly what she is obligated to do. Thus, it seems that whatever we choose to replace minimal consequentialism (i.e., C1) with should imply both that Doc is permitted to inject Pat with a proper dose of D and that Doc is permitted to inject Pat with D (though not with an improper dose of D).

3. Two Ways to Deny C1 and Solve the Problem of Act Versions

I've argued that, to solve the problem of act versions, we must reject C1 and replace it with a principle that implies that Doc is permitted to inject Pat with D even though doing so would not have optimal consequences. To do this, we have to hold that what determines whether one is permitted to perform an act like injecting Pat with D is whether there is a version of injecting him with D that has optimal consequences (such as injecting him with a proper dose of D), not whether injecting him with D would itself have optimal consequences. There seem to be only two ways to do this.

One way is to adopt maximalism and to substitute 'is morally optimal' for 'has feature F' throughout its formulation, where any option that is optimal both in terms of its consequences and in terms of all non-consequentialist factors counts as being morally optimal. And let's call the resulting view *optimalist maximalism*. And, now, recall *The Injection*. Clearly, the only morally optimal maximal options for Doc are those that entail her injecting Pat with a proper dose of D. And since these will all entail her injecting Pat with D, optimalist maximalism entails that she is permitted to inject Pat with D even though her doing so would not itself have optimal consequences. Thus, we get both the verdict that Doc's injecting Pat with a proper dose of D is permissible and that, because this entails her injecting him with D, this too is permissible, despite the fact that it wouldn't itself have optimal consequences.

The other way to get these verdicts is to adopt omnism and substitute for 'has feature F' in its formulation something that mimics optimalist maximalism's verdicts. To do so, we'll need to substitute for 'has feature F' something such as 'is entailed by a maximal option that is morally optimal' in the above formulation of omnism. And let's call this view *maximalist omnism*, because it is a version of omnism that mimics optimalist maximalism's verdicts.

The problem, though, with this way of solving the problem of act versions is that it commits the omnist to rejecting the most straightforward and plausible explanation for why Doc is permitted to perform a maximal option that is morally optimal. Rather than simply claiming, as optimalist maximalism does, that her performing such a maximal option is permissible *because it is morally optimal*, the maximalist omnist must claim that it's permissible, *because it is entailed by a maximal option that is morally optimal*. The problem is that this seems like an implausible explanation for why it's permissible to perform a maximal option that's morally optimal. Whether a maximal option is itself morally optimal seems relevant. But whether a maximal option is entailed by a maximal option that is morally optimal doesn't. Indeed, it seems no more relevant than does whether a maximal option is entailed by a maximal option that is entailed by a maximal option that is morally optimal. Thus, the second way of solving the problem of act versions forces the omnist to adopt an implausible view about what makes maximal options that are morally optimal permissible. It's not simply that they are morally optimal but rather that they are entailed by maximal options that are morally optimal.

So I've argue that, in order to solve the problem of act versions we need to adopt either optimalist maximalism or a version of omnism that, like maximalist omnism, mimics optimalist maximalism verdicts. And, of these two, we should accept the former, because it gives a more plausible and straightforward account of what makes a permissible maximal option permissible: it's not that it is entailed by a maximal option that is morally optimal or, even worse, that it is entailed by a maximal option that is entailed by a maximal option that is morally optimal, but rather simply that it is morally optimal. Thus, we should, I believe, prefer maximalism to omnism.

4. Objections to Maximalism

At this point, readers may wonder whether they must accept maximalism even if I'm right about its being preferable to 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 distinct option—e.g., the option of doing something) and non-minimal options (options that do entail some distinct option—

e.g., the option of running, which entails doing something) and adopt the following alternative to both omnism and maximalism.

Minimalism: (Min₁) For any subject S and any minimal option ω , S's ω -ing is permissible if and only if S's ω -ing has feature F, and when S's ω -ing is permissible, this is in virtue of the fact that S's ω -ing has feature F. And, (Min₂) for any subject S and any non-minimal option α , S's α -ing is permissible if and only if there exists a distinct option φ such that S is permitted to φ and S's α -ing entails S's φ -ing, and when S's α -ing is permissible, this is in virtue of the fact that there is a distinct option φ such that S's φ -ing is permissible and S's α -ing entails S's φ -ing.

But this view is a non-starter. On this view, I'm permitted to do anything that entails my doing something permissible. Thus, if I'm permitted to feed my daughter (because, say, I'm permitted to feed her a nutritious meal), then I'm also permitted to feed her shards of glass given that my feeding her shards of glass entails feeding her. And the idea that it's permissible for me to feed her shard of glass just because it's permissible for me to feed her is 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.¹² Yet I needn't do so. For my thesis is only 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. For virtually all theories of morality that have been offered to date are versions of omnism.¹³ Substitute 'maximizes utility' for 'has feature F' in the above formulation of omnism, and we get act-utilitarianism. Substitute 'accords with Kant's categorical imperative' for 'has feature F', and we get Kantianism. Substitute 'maximizes S's utility' for 'has feature F', and we get egoism.

¹² I do think that it's unlikely that there is any plausible alternative to these three. For any view besides these three must be some sort of hybrid. And that hybrid, I believe will not be superior to maximalism, for it will inherit all the problems of its non-maximalist parents. Thus, if it's part minimalism, it will have to implausibly hold that sometimes it's permissible to φ because φ -ing entails doing something that's permissible. And if it's part omnism, then it will either suffer from the problem of act versions or be forced to give an implausible explanation for why some morally optimal options are morally permissible.

¹³ As far as I know, the only exceptions are the theories defended in Brown 2015, Feldman 1986, Goldman 1978, Portmore 2011, Portmore 2013, and Zimmerman 1996, which are all versions of maximalism.

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 the objection that it cannot give a plausible solution to the problem of act versions. 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 objections. I'll consider four possible objections. Perhaps, there are other objections besides these four, but I believe that these others all fall into one of the following three categories: (1) they've already been adequately addressed elsewhere,¹⁴ (2) they're close enough to those that I do consider that they can be dismissed quickly in the footnotes below, or (3) they have not appeared in the rather small literature on maximalism and so are unknown to me.

4.1. The Professor Procrastinate Objection: Perhaps, the most significant objection to maximalism is that it is alleged to have counterintuitive implications in cases that seem to have the following three features: (2.1) S has the options of φ -ing well, φ -ing poorly, and not φ -ing at all. (2.2) It wouldn't be too bad if S were not to φ at all, but it would be much better (indeed, best) if she were to φ well. And worst of all would be if she were to φ poorly. Thus, she ought to φ well. But, (2.3) as a matter of fact, if S were to φ , she would φ poorly.

To illustrate, consider the now famous case of *Professor Procrastinate*:

Professor Procrastinate receives an invitation to review a book. He is the best person to do the review, has the time, and so on. The best thing that can happen is that he says yes, and then writes the review when the book arrives. However, suppose it is further the case that were Procrastinate to say yes, he would not in fact get around to writing the review. Not because of incapacity or outside interference or anything like that, but because he would keep on putting the task off. ...Moreover, we may suppose, [his saying yes and never writing the review] is the worst that can happen. It would lead to the book not being reviewed at all. (Jackson & Pargetter 1986, 235)

In this case, S is Professor Procrastinate, φ -ing is accepting the invitation, φ -ing well is accepting and then writing, and φ -ing poorly is accepting and then never writing.

¹⁴ One such objection to maximalism is that it generates Ross's paradox (Ross 1941). But see Castañeda 1981 and Wedgwood 2006 for why we shouldn't think that Ross's "paradox" is indeed paradoxical.

Employing this sort of case, critics offer the following argument against maximalism:¹⁵

- (2.4) Maximalism is true. [Assumption for conditional proof]
- (2.5) Thus, S's φ -ing is permissible if and only if it is entailed by some permissible option. [From 2.4 and the definition of 'maximalism']
- (2.6) S has the option of φ -ing well. [Assumption]
- (2.7) If φ -ing well is an option, it is S's best option. And, thus, if φ -ing well is an option, it's a permissible option (indeed, it is what S ought to do). [From the stipulations of the case]
- (2.8) φ -ing well entails φ -ing. [Analytic]
- (2.9) S's φ -ing is permissible. [From 2.5–2.8]
- (2.10) Thus, if maximalism is true, then S's φ -ing is permissible. [From 2.4–2.9 by conditional proof]
- (2.11) It is not the case that S's φ -ing is permissible. [Intuition]
- (2.12) Therefore, it is not the case that maximalism is true. [From 2.10 & 2.11]

Such critics find 2.11 intuitively obvious. They claim that Professor Procrastinate is not permitted to accept the invitation given that he would not write if he were to do so.

Of course, not everyone finds 2.11 intuitively compelling, but that's the argument that many critics give. I suspect that disagreement about 2.11 stems from the fact that such cases (e.g., *Professor Procrastinate*) are under-describe. And, once we take note of the two

¹⁵ See, for instance, Cariani 2013, Jackson & Pargetter 1986, and Snedegar 2014. Admittedly, they don't use this sort of argument directly against maximalism. Instead, they use it against similar views, such as possibilism (i.e., the view that whether a subject ought to φ depends on whether the best course of action that's possible for her entails her φ -ing) and the principle of inheritance (i.e., the view that if φ -ing entails ψ -ing, then 'S ought to φ ' entails 'S ought to ψ '). But, given the relevant similarities between these views and maximalism, the argument can be turned into one against maximalism, as I've done above.

As I see it, possibilism is just a specific version of maximalism: one that, unlike maximalism, is committed to a certain account of what our options are. For possibilists seem to assume something like the following account: S has, as of t , the option of φ -ing at t' ($t < t'$) so long as there is some available schedule of intentions beginning at t such that, if S's intentions were to follow this schedule, S would φ at t' . For examples of such possibilists, see Feldman 1986 and Zimmerman 1996. For criticisms of this account of options, see Portmore 2011 and Portmore Forthcoming.

Maximalism is also broader than the principle of inheritance. Although maximalism entails this principle, this principle doesn't entail maximalism. This is because the principle of inheritance doesn't contain any explanatory component, nor does it say anything about when, if ever, S ought to perform a maximal option (i.e., an option that is not entailed by any other option).

ways in which the missing details might be spelled out, we see that we should reject either 2.6 or 2.11. Here's my argument for this.

- (2.13) S is either irrepressible or repressible — that is, either (a) S will φ poorly regardless of how he now responds to his reasons or (b) it is not the case that S will φ poorly regardless of how he now responds to his reasons. [From the law of excluded middle]
- (2.14) If S is irrepressible, then we should reject 2.6, which says that S has the option of φ -ing well. For if S will φ poorly regardless how he now responds to his reasons, then he doesn't, at present, have the power to direct the course of his future actions such that he will φ well. And, thus, he doesn't, at present, have the option of φ -ing well.¹⁶ [Intuition]
- (2.15) If S is repressible, then we should reject 2.11, which says that S's φ -ing is not permissible. For, if he is repressible, then he should direct the course of his future actions by responding appropriately, at present, to his reasons, thereby ensuring that he will φ and φ well. And, thus, he's not just permitted to φ ; he's obligated to φ . [Intuition]
- (2.16) Therefore, we should reject either 2.6 or 2.11. [From 2.13–2.15]

Note that it's unclear from Jackson and Pargetter's description of *Professor Procrastinate* whether Professor Procrastinate is repressible or irrepressible. For all that they say, then, it could be that Procrastinate is aware of his tendency to procrastinate and that, when it's really important to him that he doesn't procrastinate, he resolves now not to give into the temptation to procrastinate later on. And it may even be, as we'll indeed suppose, that his making this resolution now is sufficient to ensure that he won't procrastinate later on. And, in that case, he is repressible, for he will write the review if he responds

¹⁶ I suspect that possibilists will balk at this claim. For, as pointed out in the previous note, they hold that S has, as of t , the option of φ -ing at t' ($t < t'$) so long as there is some available schedule of intentions beginning at t such that, if S's intentions were to follow this schedule, S would φ at t' . Thus, they deny that for S to have, as of t , the option of φ -ing at t' , there has to be some way for her to respond at t to her reasons such that this would result in her φ -ing at t' ($t < t'$). But my hope is only to convince those who do not have any prior commitment to possibilism, and, importantly, this includes people who, like Jackson and Pargetter, offer the 2.4–2.16-type argument to which I'm responding. And, for both a critique of the possibilist's account of options (which I call *schedulism*) and a defense of an account that vindicates 2.14, see Portmore Forthcoming.

appropriately to his reasons by resolving now, as he accepts the invitation, to write the review as soon as the book arrives. So, one possibility for why he wouldn't write if he were to accept is that he's not now responding appropriately to his reasons. And, in that case, my clear intuition is that Professor Procrastinate is not only permitted to accept, but is obligated to accept. For he's obligated to respond appropriately to his reasons, accepting the invitation while also resolving to write the review as soon as the book arrives. In which case, he will accept and write the review. And so it seems that we should reject 2.11 if he is repressible.¹⁷

Of course, he could instead be irrepressible such that no matter how he responds now to his reasons, and, thus, no matter what he resolves now to do later on, his later self is going to choose to procrastinate when the book arrives. In that case, I think that he just has to accept that he is irrepressible, having no more control at present over whether his future self will write the review than I have over whether the next U.S. Congress will amend the constitution so as to prohibit the private ownership of firearms. And, if that's how we're supposed to imagine things, then although we should readily accept that Professor Procrastinate is not permitted to accept (and, so, concede that 2.11 is true), we should deny that he has, at present, the option of accepting and writing—that is, deny 2.6. For if he doesn't, at present, have the power to direct the course of his future actions so as to ensure that he will write the review when the book arrives, in what sense does he have, at present, the option of writing when the book arrives? To have the option of ϕ -ing is to have control over whether one ϕ s.¹⁸ But, if Professor Procrastinate is irrepressible, he has, at present, no control over whether he will write when the book arrives. And so I think that we should deny 2.6 if Procrastinate is irrepressible.

¹⁷ Another possibility is that he has made the following prior arrangement with a colleague: If he copies this colleague on an email in which he accepts an invitation to do something, he thereby bets this colleague ten thousand dollars that he will do that something. Further suppose that, given this, Professor Procrastinate would write the review if he were to copy this colleague on the email in which he accepts the journal's invitation to write the book review. And this is compatible with Jackson and Pargetter's description of the case, for it may be that Professor Procrastinate doesn't want to make this bet and so wouldn't copy this colleague on the email if he were to accept. And, in that case, it could still be that, were Professor Procrastinate to accept, he would not write the review. But, in such a case, I do not have the intuition that Professor Procrastinate should not accept the invitation (that is, I do not find 2.11 plausible). Rather, I have the intuition that Professor Procrastinate should accept by email while copying this colleague on that email and should, therefore, accept. Even Frank Jackson (2014, pp. 645–646) now rejects 2.11 in cases like this that involve simultaneous actions, such as where emailing the journal to accept the invitation and copying this colleague on that email occur simultaneously.

¹⁸ See Portmore Forthcoming for further defense of this claim.

4.2. *Gustafsson's Objection*: Another recent objection to maximalism comes from Johan Gustafsson.¹⁹ He argues that maximalism has counterintuitive implications in the following sort of case.

Newcomb's Non-Problem: At t_0 , you are offered a chance to participate in a non-problematic version of Newcomb's problem, which involves two boxes: one transparent and one opaque. At t_1 , you must choose whether to participate. If you agree at t_1 to participate, you must take at t_3 either just the transparent box or both boxes. And, as you are told, you must take possession of the contents of whatever box or boxes you take, whether they be good or bad. Thus, if you take both boxes, and the opaque box happens to contain a writ of debt for a \$1,000,000, you will be responsible for paying this debt. Now, as you can see, the transparent box contains \$1,000. But the contents of the opaque box are a complete mystery to you — that is, until t_2 , which is when you are given the following additional information: the opaque box contains either \$1 or \$1,000,001. It contains \$1 if you formed at t_1 the intention to take both boxes. Otherwise, it contains \$1,000,001.

Even if you end up choosing at t_1 to participate and taking at t_3 both boxes, there are two possibilities, which have very different outcomes:

(Poss1) At t_1 , you chose to participate. At t_3 , you took both boxes. But you did not form the intention to take both boxes until t_2 , which is when you learned that there's nothing bad in the opaque box. Consequently, you ended up with \$1,001,001.

(Poss2) At t_1 , you chose to participate. And, at t_1 , you formed the intention to take both boxes at t_3 even though you did not at the time have any idea what was in the opaque box. At t_3 , you took both boxes. You ended up with only \$1,001.

¹⁹ See his 2014. I've revised his case so that even those who deny that there are pragmatic reasons for forming intentions (thus, denying that the fact that intending to φ , which is distinct from φ -ing, would have good consequences is a reason to intend to φ) will still accept Gustafsson conclusions about which intentions the agent should form.

Clearly, if Poss2 is actualized, you have failed in some way. For, as Gustafsson notes, it was up to you at t_0 which of Poss1 and Poss2 would be actualized. And if Poss2 is actualized, then you ended up with a million fewer dollars. Of course, from the mere fact that you ended up with less money, it doesn't necessarily follow that you failed in some way. But the thought is that you ended up with less money than you should have ended up with, which does imply that you failed in some way. The thought is that you should not have formed at t_1 the intention to take both boxes at t_3 . That was reckless given that, for all you knew, the opaque box contained a writ of debt for a \$1,000,000 or more. Moreover, had you not recklessly formed this intention at t_1 , you would have actualized Poss1, ending up with an additional million dollars. Thus, as Gustafsson claims, it seems that a plausible theory would require you to actualize Poss1 as opposed to Poss2. The problem with maximalism, Gustafsson believes, is that maximalism cannot do this. But while it's true that some versions of maximalism cannot do this, there is a plausible version of maximalism that can.

On exclusively act-orientated versions of maximalism, only voluntary actions constitute options. Thus, your only options are (O1) choose at t_1 to participate and take at t_3 both boxes, (O2) choose at t_1 to participate and take at t_3 only the transparent box, and (O3) choose at t_1 not to participate. It takes these to be your only options, because only that which can be done at will counts as an option, and your forming or refraining from forming an intention is not something that you can do at will—that is, you do not form or refrain from forming an intention by intending to do so.²⁰ So the best that exclusively act-orientated versions of maximalism can do is require you to perform O1, but your performing O1 doesn't necessitate your actualizing Poss1 as opposed to Poss2. Thus, there is no way for such versions of maximalism to require you to actualize Poss1 as opposed to Poss2.

²⁰ To see that the formation of an intention is not something one can do at will, consider the toxin puzzle (Kavka 1983). I will receive a million dollars tomorrow morning if and only if, at midnight tonight, I intend to drink some toxin tomorrow afternoon. Drinking the toxin will not kill me, but it will make me terribly ill for several days. Whether I receive the million dollars tomorrow morning depends only on what I intend to do at midnight tonight, not on whether I drink the toxin tomorrow afternoon. Realizing this, I'm unable to intend at midnight tonight to drink the toxin tomorrow afternoon. For I see no reason to drink the toxin. I know that, come tomorrow afternoon, I'll either have the million dollars or I won't. And in neither case will I have any reason to drink toxin. So, I have decisive reason not to drink the toxin. Given this, I'm unable to form the intention to drink the toxin. And this shows that forming an intention is not typically something that I can do at will—that is, by forming the intention to form that intention.

Fortunately, there is no reason for us to adopt an exclusively act-orientated version of maximalism. Intuitively, it seems that we have obligations and responsibilities with respect to more than just voluntary acts. For instance, it seems that we have obligations and responsibilities with respect to such things as feeling contrite, not forgetting our anniversary, believing in accord with our evidence, wanting what's best for our children, intending to take what we judge to be the necessary means to our ends, etc. Yet none of these are things that we do at will. Given this, we should not restrict maximalism to only voluntary acts.²¹ We should instead think that maximal options will include both voluntary acts as well as the formation of various reasons-responsive attitudes, such as beliefs, desires, and intentions. And when we do, we see that the relevant options are (O2) choose at t_1 to participate and take at t_3 only the transparent box, (O3) choose at t_1 not to participate, (O4) choose at t_1 to participate while simultaneously forming the intention to take both boxes and then take at t_3 both boxes, and (O5) choose at t_1 to participate while not forming at t_1 the intention to take both boxes, form at t_2 the intention to take both boxes, and take at t_3 both boxes. Since maximalism so construed allows that O5 is an option, and since your performing O5 necessitates your actualizing Poss1, maximalism can require you to actualize Poss1 by requiring you to perform O5. Thus, maximalism can avoid Gustafsson's objection.²² It's just that Gustafsson wrongly assumed, and without argument, that maximalism must be exclusively act-orientated.²³

4.3. *The 'Actual'-Implies-'Option' Objection:* Some may worry that maximalism has counterintuitive implications in the following sort of case.

Teeing Off: Alejandro is playing golf and is about to tee off on the second hole when an evil demon presents him with the following dilemma. If he quits the game,

²¹ For more on this, see Portmore Forthcoming.

²² See Bykvist 2002 for another objection to maximalism, and see Gustafsson 2014 (587–588) for some discussion. As with the objections discussed in the body of this paper, it relies on an implausible account of what our options are. For Bykvist's objection to even get off the ground, he must presume that S has, as of t , the option of φ -ing at t' only if she can form at t the intention to φ at t' . On this view, I don't now have the option of forming tomorrow the belief that there are two sticks of butter in my fridge (assume that tomorrow there will be exactly two sticks of butter in my fridge) even though I can now form the intention to look inside my fridge tomorrow and would form this belief if I were to do so. On Bykvist's view, I don't now have the option of forming this belief tomorrow, because I cannot now form the intention to form this belief. For I'm a rational person, and a rational person cannot form the belief to φ when she knows that φ is not something that she can intentionally do.

²³ Regan 1980 and Portmore Forthcoming for why we should not accept any theory that is exclusively act-orientated.

refusing to tee off, the evil demon will do nothing. If he tees off and slices the ball, the evil demon will kill everyone. If he tees off and doesn't slice the ball, the evil demon will give everyone a tiny reward. Now, although Alejandro can easily control whether he tees off, he doesn't, when teeing off, have control over whether he slices the ball or not. About half the time, he slices the ball whether he tries to or not.

Now, it may seem that Alejandro has all the following options: (Opt_i) teeing off; (Opt_{ii}) not teeing off; (Opt_{iii}) teeing off and slicing the ball; and (Opt_{iv}) teeing off and not slicing the ball. And if this is right, then it would seem both that (1) Alejandro's teeing off (that is, Opt_i) is impermissible given how risky it would be and that (2) Alejandro is permitted to tee off and not slice the ball (that is, Opt_{iv}) given that this would be better than any alternative option. Yet these two verdicts are incompatible with maximalism, for maximalism entails that, if teeing off and not slicing is permissible, then teeing off must also be permissible given that teeing off and not slicing entails teeing off (that is, Opt_{iv} entails Opt_i).²⁴

But I think that we should deny that teeing off and not slicing (that is, Opt_{iv}) is an option for Alejandro in the relevant sense—that is, in the sense that's relevant to determining whether his teeing off and not slicing is the sort of thing that he could be obligated to do and accountable for failing to do. For teeing off and not slicing to be an option for him in this sense, it seems that it would have to be something that he has control over.²⁵ And it seems that although he has control over both whether he tees off and whether he tries to do so without slicing, he does not have control over whether he slices the ball when he tees off. Of course, suppose that he in fact wrongly decides to tee off and just so happens not to slice the ball. Someone might insist that if he actually tees off and not slices the ball, it follows that he must have had the option of teeing off and not slicing the ball. The assumption, then, is that 'S actually φ -ed' implies that 'S had the option of φ -ing'. But, as others have argued, I think that we should deny this; from the fact that S actually φ -ed, it doesn't follow that S ever had the relevant sort of control over whether she was to φ .²⁶

²⁴ Thanks to Campbell Brown for pressing me on this.

²⁵ For more on this, see Portmore Forthcoming.

²⁶ For more on this, see Southwood and Wiens Forthcoming.

So I would argue that his options are: (Opt_I) teeing off, (Opt_{II}) not teeing off, (Opt_V) teeing off while trying to slice the ball, and (Opt_{VI}) teeing off while trying not to slice the ball. And, so although Opt_{VI} entails Opt_I, this is no problem for maximalism, for neither is permissible. They are both just too risky. So, when we take note of what Alejandro's options actually are, we find that cases such as *Teeing off* present no difficulty for maximalism.

4.4. *The Implausible Explanation Objection*: So far, I've dealt with objections that allege that maximalism has counterintuitive implications. In each case, I've shown that maximalism gets counterintuitive implications only if combined with an implausible account of what our options are. And so these objections are not actually objections to maximalism per se, but only to certain implausible substantive versions of maximalism. But there is an entirely different sort of objection to maximalism that I need to address. For one might argue that the problem with it is not that it gets the wrong verdicts but that it provides an implausible account of those verdicts.²⁷ Let's consider this objection.

According to maximalism, if a subject is permitted to perform a non-maximal option *v*, this is in virtue of the fact that her *v*-ing is entailed by a maximal option that has feature *F*. But it may seem that the explanation for why she is permitted to *v* lies simply with the fact that her *v*-ing is itself *F* (*F* being some fundamental right-making feature) and not with the fact that it is entailed by a maximal option that is *F*. To illustrate, suppose that I'm getting ready for bed and that, under the circumstances, it is permissible for me to brush my teeth as part of my preparations. Now, the explanation for this being permissible may seem to lie simply with the fact that my doing so would itself have optimal consequences (by preventing tooth decay) and not with the fact that it is entailed by a maximal option that would have optimal consequences. So it may seem that, contrary to maximalism, maximal options have nothing to do with the explanation for why I'm permitted to brush my teeth.

Yet, if I'm permitted to brush my teeth, it must be in virtue of my being permitted to perform some specific instance of teeth-brushing. After all, if my circumstances were different such that the only brush available to me was a wire-bristled one, I wouldn't be permitted to brush my teeth. For brushing my teeth with such a brush would cause more damage to my teeth than not brushing at all. Also, if my circumstances were such that there

²⁷ Thanks to Shyam Nair for pressing me on this.

was no permissible way for me to follow up with brushing my teeth, I wouldn't be permitted to brush my teeth. Suppose, for instance, that there was an evil demon who would kill me and my family if I were to brush my teeth without following up by drinking a glass of orange juice immediately afterwards. In that case, I would not be permitted to brush my teeth. For it would be bad to brush my teeth and then drink a glass of orange juice immediately afterwards. The combination of brushing and then immediately drinking something acidic would cause more damage to my teeth's enamel than not brushing at all. And I certainly wouldn't be permitted to brush my teeth and not follow up by drinking the orange juice immediately afterwards given that this would result in the evil demon's killing me and my family.

So it seems that if I'm permitted to brush my teeth, it must be because I'm permitted to perform some option that entails my brushing my teeth (e.g., the option of brushing my teeth with a toothbrush in the dentist-recommended fashion and then following up by going straight to bed without eating or drinking anything). And, in turn, if I'm permitted to perform this option, it must be because I'm permitted to perform some option that entails performing this option. And so and so forth, until we reach a maximal option that is F (e.g., has optimal consequences). So, on reflection, we should think that the explanation for why I'm permitted to brush my teeth has everything to do with its being entailed by a maximal option that is F.

To further illustrate the plausibility of maximalism's explanatory component, consider the following case.

Two Drugs: Dr. Singh is deliberating at t_0 about what drugs, if any, to give her patient, Prasad, at t_2 . In fact, there are exactly two drugs available to her: A and B. If she gives Prasad 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 one of A and B

at t_2 . Moreover, she's going to follow through with this intention. For, as it turns out, she has a grudge against Prasad. Given that she's going to form, and then follow through, with the intention to give him just one of A and B at t_2 , the following two subjunctive conditionals are true: (SC₁) if she were to give him A at t_2 , she would not give him B at t_2 and (SC₂) if she were to give him B at t_2 , she would not give him A at t_2 . Nevertheless, if she were to respond appropriately to the decisive 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 everyone would live happily ever after. Lastly, assume that Dr. Singh will lose her medical license if she fails to cure him and will go to prison if she kills him.

Now, it seems that Dr. Singh is both permitted to give Prasad A at t_2 and permitted to give him B at t_2 . After all, she's permitted (indeed, obligated) to give him both A and B at t_2 , and this entails both giving A at t_2 and giving him B at t_2 . But the explanation for why each of these acts are permissible has nothing to do with the facts about each. For instance, the explanation has nothing to do with the consequences of each of these acts. For, given the truth of SC₁, injecting him with A at t_2 would have terrible consequences. And, given the truth of SC₂, injecting him with B at t_2 would also have terrible consequences. And it can't be that what explains the permissibility of each of these acts is that each would have terrible consequences. Instead, the explanation for why Dr. Singh is both permitted to give him A at t_2 and permitted to give him B at t_2 must lie with the fact that she's permitted to do something that entails doing each of these—something that would have good consequences. After all, she is permitted to give Prasad both A and B at t_2 , and this entails both giving him A at t_2 and giving him B at t_2 . Taking this further, it would seem that the explanation for why she is permitted give Prasad both A and B at t_2 is that she is permitted to perform some option that entails doing this, such as injecting him with these two drugs using sterile needles. And so and so forth, until we reach some maximal option that is F.

So although it may seem strange at first to think that the explanation for its being permissible to perform some non-maximal option lies with the fact that it is entailed by the performance of a maximal option that is F, these cases illustrate that, on reflection, this is indeed what we should think. We should think, for instance, that it is permissible for Dr.

Singh to inject Prasad with A at t_2 , not because doing so would itself have optimal consequences (it wouldn't), but because it is entailed by a maximal option that would have optimal consequences—the one in which she injects Prasad with both A and B at t_2 .

5. Conclusion

I've argued that we should prefer maximalism to omnism. Maximalism stems from the intuitive idea that, if I'm permitted to perform a certain type of action, this is because there is some instance of this type that I'm permitted to perform. Moreover, I've shown both that maximalism offers a plausible solution to the problem of act versions and that none of the objections that have typically levelled against it succeed. By contrast, omnism is unable to offer a plausible solution to the problem of act versions.

If we should favor maximalism over omnism, then nearly all extant theories of morality are in need of substantial revision, for nearly all of them are versions of omnism.²⁸ Of course, having to reject all the many versions of omnism that have been offered to date doesn't mean that all is lost. We just need to reformulate these theories. These theories identify some right-making feature of maximal actions, and, for all that I've said, they may have correctly done so. They just make the mistake of thinking that we are to assess all options in terms of this feature when we must assess only maximal options in terms of this feature. The lesson, then, is that Kantians, utilitarians, Rossian pluralists, and others need to revise their theories on the model of maximalism, inserting what they take to be the relevant right-making feature for F in that formulation. Still, it's important for these theorists to reformulate their views as versions of maximalism. For, otherwise, as I've shown, their theories will suffer either from the problem of act versions or from the problem of giving an implausible account of what makes permissible maximal options permissible. So, they may have figured out what F is, but, if I'm right, they've misunderstood which options are to be evaluated directly in terms of F and which are instead to be evaluated in terms of options that entail them.²⁹

²⁸ For a list of the exceptions, see note 13 above.

²⁹ For helpful comments and questions, I thank Brad Armendt, Elizabeth Brake, Campbell Brown, Brad Cokelet, Seth Lazar, Thomas Mautner, David McElhoses, N. Ángel Pinillos, Felix Pinkert, Steven Reynolds, and audiences at Australian National University and Charles Sturt University, Wagga Wagga. I especially thank Wylie Breckenridge, Daniel Cohen, and G. Shyam Nair for some extraordinarily thorough and helpful comments. Work on this paper was supported by the RSSS Visiting Fellows Program, School of Philosophy, Australian National University.

Arizona State University
 dwportmore@gmail.com

References

- Bergström, L. (1966). *The Alternatives and Consequences of Actions*. Stockholm: Almqvist & Wiksell.
- 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.
- Cariani, F. (2013). "'Ought' and Resolution Semantics." *Noûs* 47: 534–558.
- Castañeda, H.-N. (1968). "A Problem for Utilitarianism." *Analysis* 28: 141–142.
- . (1981). "The Paradoxes of Deontic Logic: The Simplest Solution to All of Them in One Fell Swoop." In R. Hilpinen (ed.), *New Studies in Deontic Logic*, pp. 37–85. Dordrecht: D. Reidel.
- Feldman, F. (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht: D. Reidel Publishing Company.
- Goldman, H. S. [now H. M. Smith]. (1978). "Doing the Best One Can." In A. I. Goldman and J. Kim (eds.), *Values and Morals*, pp. 185–214. Dordrecht: D. Reidel Publishing Company.
- Gustafsson, J. E. (2014). "Combinative Consequentialism and the Problem of Act Versions." *Philosophical Studies* 167: 585–596.
- Jackson, F. (2014). "Procrastinate Revisited." *Pacific Philosophical Quarterly* 95: 634–647.
- Jackson, F. and R. Pargetter (1986). "Oughts, Actions, and Actualism." *Philosophical Review* 95: 233–255.
- Kavka, G. (1983). "The Toxin Puzzle." *Analysis* 43: 33–36.
- Kiesewetter, B. (2015). "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125: 921–946.
- Portmore, D. W. (Forthcoming). "Maximalism and Moral Harmony." *Philosophy and Phenomenological Research*.
- . (2016). "Maximalism vs. Omnism about Reasons." Working manuscript that's available at <http://bit.ly/1Mygc3m>.
- . (2013). "Perform Your Best Option." *The Journal of Philosophy* 110: 436–459.

- . (2011). *Commonsense Consequentialism: Wherein morality meets rationality*. New York: Oxford University Press.
- Regan, D. (1980). *Utilitarianism and Co-operation*. New York: Oxford University Press.
- Ross, A. (1941) "Imperatives and Logic." *Theoria* 7: 53–71.
- Snedegar, J. (2014). "Deontic Reasoning Across Contexts." In Cariani, F., Grossi, D., Meheus, J., and Parent, X. (eds.), *Deontic Logic and Normative Systems*, pp. 208–223. Switzerland: Springer International Publishing.
- Southwood, N. and D. Wiens. (Forthcoming) "'Actual' Does Not Imply 'Feasible'." *Philosophical Studies*.
- Wedgwood, R. (2006). "The Meaning of 'Ought'." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Volume 1, pp. 127–160. Oxford: Oxford University Press.
- Zimmerman, M. J. (1996). *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.