

Morality, Rationality, and Performance Entailment*

Douglas W. Portmore

Abstract: The performance of one option can entail the performance of another. For instance, baking an apple pie entails baking a pie. Now, suppose that both of these options—baking a pie and baking an apple pie—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 bake an apple pie? Or is baking an apple 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 accord with Kant's categorical imperative? I defend the view that the permissibility of an option that entails another is more fundamental than the permissibility of the option that it entails. That is, I defend *maximalism*: the view that if an agent is permitted to perform a certain type of action (say, baking a pie), this is in virtue of the fact that she is permitted to perform some instance of this type (say, baking an apple pie), where φ -ing is an instance of ψ -ing if and only if φ -ing entails ψ -ing but not vice versa. If maximalism is correct, then, as I show, most theories of morality and rationality must be revised.

Keywords: maximalism, morality, obligations, options, Professor Procrastinate, rationality.

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 typing and the option of typing the word 'the', and my typing the word 'the' entails my typing, for it is logically impossible for me to type the word 'the' without typing. Such instances of performance entailment are ubiquitous. Kissing passionately entails kissing. Walking while chewing gum entails walking. Driving under 55 mph entails driving under 100 mph. Marrying an unmarried man entails marrying a bachelor. Stretching a t_1 and then going for a run at t_2 entails going for a run at t_2 .

Given that our options include both those that entail others and those that are entailed by others, it seems that our practical normative theories (e.g., our theories of morality and rationality) need to be able to assess both types. For instance, if I have the options of (Opt1) baking, (Opt2) baking a pie, (Opt3) baking an apple pie, and (Opt4) baking an apple pie and then taking it to the family who just moved in across the street, it won't do for our normative theory to tell us only whether some proper subset of these are permissible. If these are all equally options, then a normative theory concerning the permissibility of an agent's options owes us an account of the permissibility of each and every one.

Now, suppose that our normative theory tells us that all of Opt1–Opt4 are permissible. This raises the following important question: What is the relationship between the permissibility of an option that entails another (e.g., Opt3) and the permissibility of the option that it entails (e.g., Opt2)? More specifically, which, if either, is more fundamental: the permissibility of the option that entails the other or the permissibility of the option that it entails? Is baking a pie permissible because it's permissible to bake an apple pie? Or is baking an apple 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 accord with Kant's categorical imperative? Of course, it could be that sometimes the option that entails is more fundamental, but other times the option that is entailed is more fundamental. There are, then, exactly four mutually exclusive and jointly exhaustive possibilities: (Poss1) the permissibility of an option that entails another is always *more* fundamental than the permissibility of the option that it entails; (Poss2) the permissibility of an option that entails another is always *less* fundamental than the permissibility of the option that it entails; (Poss3) the permissibility of an option that entails another and the permissibility of the option that it entails are always *equally* fundamental—that is, the one is *never either more or less* fundamental than the other; and (Poss4) there is a lack of uniformity, as there would be, for instance, if sometimes the permissibility of an option that entails another is more fundamental, but other times it is less fundamental.

The point of this paper is to defend Poss1. I'll defend it against some typical objections. And I'll argue that Poss1 is more plausible than its three alternatives. I'll proceed by considering each of the four possibilities in turn, discussing their relative merits and demerits as we go.

1. Maximalism

The first possibility, Poss1, holds that the permissibility of an option that entails another is always more fundamental than the permissibility of the option that it entails. This view is sometimes called *maximalism*.¹ On this view, we must distinguish between maximal and

¹ See, for instance, Bykvist 2002 and Gustafsson 2014. Note, however, that my use of the term is broader than theirs. Whereas they use the term to refer only to *consequentialist theories* that take the *moral* permissibility of maximal options to be more fundamental than that of non-maximal options, I use the term to refer to *any theory* that takes the (*moral or rational*) permissibility of maximal options to be more fundamental than that of non-maximal options.

non-maximal options. An option φ is a *maximal option* if and only if there is no option ψ such that ψ -ing entails φ -ing but not vice versa. Otherwise, it's a *non-maximal option*. So walking at t_1 won't be a maximal option if walking fast at t_1 is an option. And walking fast at t_1 won't be a maximal option if walking fast at t_1 and then jogging at t_2 is an option. And so on and so forth until we get to an option for which there is no other option that entails it.

We must distinguish between these two types of options, because, on maximalism, they are to be assessed differently. For, when fully spelled out, the view says the following.

Maximalism: (Max1) For any non-maximal option ν , S 's ν -ing is permissible if and only if there exists an option φ such that S 's φ -ing is permissible and S 's φ -ing entails S 's ν -ing but not vice versa, and when S 's ν -ing is permissible, this is in virtue of the fact that there exists an option φ such that S 's φ -ing is permissible and S 's φ -ing entails S 's ν -ing but not vice versa. And, (Max2) for 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 .²

Maximalism is neutral on what sorts of things count as options. Perhaps, it is only voluntary acts. But, perhaps, the non-voluntary formations of beliefs and other attitudes also count as options. And maximalism allows for options to be conjunctive. An agent could, for instance, have the option of asserting that p while believing that p is false. What's more, maximalism is neutral on what ' F ' stands for. Indeed, we can substitute for 'has feature F ' anything that would render Max2 coherent, including 'maximizes utility', 'accords with Kant's categorical imperative', or 'contains only beliefs for which S has sufficient evidence and contains all and only those acts that maximize expected utility'. As this last example illustrates, F can be something quite complicated, involving the assessment of different types of options according to different criteria—e.g., evaluating beliefs in terms of evidence and acts in terms of expected utility.

² There are two types of maximalism: moral maximalism and rational maximalism. To get moral maximalism, insert the word 'morally' before each instance of 'permissible' in the above formulation. To get rational maximalism, insert the word 'rationally' instead. In this paper, I will be concerned with both moral maximalism and rational maximalism and so will typically leave the kind of permissibility (moral or rational) unspecified. To facilitate this, I will discuss only examples in which the agent's moral permissions coincide with her rational permissions.

The thought underlying maximalism is that if I'm permitted to perform a certain type of action (say, baking), then there must be some instance of this type that I'm permitted to perform (say, baking a pie). For if I'm not permitted to bake anything (not a pie, not a cake, not cookies, not anything), then I'm not permitted to bake. Moreover, if I'm not permitted to bake a pie and then eat either some or none of it (and assume that, if I bake a pie, I must then eat either some or none of it), then I'm not permitted to bake a pie. Suppose, for instance, that an evil demon will kill me and everyone I love if either I bake a pie and then eat some of it or bake a pie and then eat none of it. In that case, it would be impermissible for me to bake a pie. And this is because there is no instance of baking a pie that I'm permitted to perform. Note, then, that, as I'm understanding things, S 's φ -ing counts an instance of S 's ψ -ing if and only if S 's φ -ing entails S 's ψ -ing but not vice versa. Thus, both baking-a-pie-and-then-eating-some-of-it and baking-a-pie-and-then-eating-none-of-it count as instances of baking a pie.

So, 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 an apple pie. And if I'm permitted to bake an apple pie, this is because I'm permitted to perform some instance of apple-pie-baking, such as baking an apple pie and then taking it to the family who just moved in across the street. And so on and so forth. But, of course, if this just goes on forever, we'll end up with an infinite regress. I believe, however, that this will not go on forever. We will eventually arrive at an option that is not entailed by any other option.

Some may question this belief. The worry would be that, for any option (e.g., the option of punching), there will always be a more specific option that entails it (e.g., the option of punching a speed bag) and an even more specific option that entails that option (e.g., the option of punching a speed bag softly), and so on and so forth, *ad infinitum*. It may seem, then, that we will never arrive at an option that is so specific that there is no other more specific option that entails it.

But although there is no limit to the degree of specificity with which we can describe a particular action, there is a limit to an agent's ability to determine the specificity of her actions. Consider that although I can determine whether I punch a speed bag with my right or left fist and also whether I punch it softly or forcefully, I can't, it seems, determine

whether I punch it with precisely 113.1052 newtons of force. And, importantly, I must be able to control whether I ϕ if my ϕ -ing is to count as an option for me and, thus, as something that I can be obligated to perform and be accountable for failing to perform. But, in the relevant sense, I do not control whether I punch the bag with precisely 113.1052 newtons of force, and so this does not count as an option for me. And this explains both why I cannot be obligated to do so and why I can't be held accountable should I fail to do so. So even though I have the option of punching the bag softly and my punching the bag with 113.1052 newtons of force entails punching the bag softly, I don't have the option of punching it with precisely 113.1052 newtons of force. For the control that I exert over the specificity of my actions is limited. Consequently, the specificity of my options is limited. Thus, it is not the case that, for any option that I have, I must have some other more specific option that entails it. Although there will always be a more specific event that entails it, that more specific event will not always be an option for me. And this along with the fact that our existences are finite suggests that we will always eventually arrive at a maximal option.

When we do arrive at a maximal option, we won't be able to derive its permissibility from that of some other permissible option that entails it. Given that it's a maximal option, there will be no other (distinct) option that entails it. Thus, maximalism must include not only Max1 but also Max2. Max2 tells us that a maximal option is permissible if and only if it has feature F. Moreover, Max2 tells us that when a maximal option is permissible, this is in virtue of the fact that it has feature F. But what's distinctive and interesting about maximalism is that it tells us that we don't evaluate non-maximal options in the same way. Instead of evaluating non-maximal options in terms of whether or not they have feature F, we evaluate them in terms of whether there is some other permissible option that entails it. Thus, Max1 tells us that a non-maximal option is permissible if and only if there is some permissible option that entails it but not vice versa. Moreover, it says that when a non-maximal option is permissible, this is in virtue of the fact that there is some permissible option that entails it *but not vice versa*. The need for this italicized phrase can be seen when we consider pairs of options that each entail the other, such as marrying a bachelor and marrying an unmarried man. It can't be that the one is permissible in virtue of being entailed by the other. For to say that an option is permissible in virtue of being entailed by itself (or its logical equivalent) is to offer no explanation at all. What explains the

permissibility of marrying an unmarried man, then, is not that it is permissible to marry a bachelor, but that it is permissible to marry some specific bachelor (i.e., some specific unmarried man).

So, at this point, the reader should have a basic understanding of maximalism. This understanding will be deepened as we consider its implications and compare it to its competitors. But before we look to some of maximalism's competitors, it will be good to consider some of the standard objections to maximalism and why I think they fail.

2. Objections to Maximalism

Perhaps, the most significant of these is the objection that maximalism has 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 would be okay if S doesn't φ at all, but it would be 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. (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.

Employing this sort of case, critics offer the following argument against maximalism:³

³ 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 views like possibilism (i.e., the view that whether one ought to φ depends on whether the best possible thing that one can do entails φ -ing) and inheritance (i.e., the view that if φ -ing

- (2.4) Maximalism is true, and, thus, S's φ -ing is permissible if and only if it is entailed by some permissible option. [Assumption for *reductio*]
- (2.5) S has the option of φ -ing well. [Assumption]
- (2.6) 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]
- (2.7) φ -ing well entails φ -ing. [Analytic]
- (2.8) Thus, S's φ -ing is permissible. [From 2.4–2.7]
- (2.9) S's φ -ing is not permissible. [Intuition]
- (2.10) Therefore, it is not the case that maximalism is true. [From 2.4, 2.8, & 2.9]

Such critics find 2.9 intuitively obvious. They claim that Professor Procrastinate is not permitted to accept given that he would not write if he were to accept.

Now, not everyone finds 2.9 intuitively compelling, but nevertheless that's the argument against maximalism that many critics give. I suspect that disagreement about 2.9 stems from the fact that such cases (e.g., *Professor Procrastinate*) are under-describe. And, once we take note of the two ways in which the missing details might be spelled out, we see that we should reject either 2.5 or 2.9. Here's my argument for this.

- (2.11) 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]
- (2.12) If S is irrepressible, then we should reject 2.5, 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. [Intuition]

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.

- (2.13) If S is repressible, then we should reject 2.9, 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]
- (2.14) Therefore, we should reject either 2.5 or 2.9. [From 2.11–2.13]

To illustrate, consider *Professor Procrastinate*. Given Jackson and Pargetter's description of the case, it's unclear whether Professor Procrastinate is repressible or irrepressible. For all that they say, 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 is sufficient to ensure that he won't procrastinate. And, in that case, he is repressible, for he will write the review if he responds appropriately to his reasons by resolving now 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 2.9 if he is repressible.⁴

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

⁴ Another possibility is that he has made the following prior arrangement with a colleague: If he copies his 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 his 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 his 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 2.9 plausible). Rather, I have the intuition that Professor Procrastinate should accept by email while copying his colleague on that email and should, therefore, accept.

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, accept 2.9), we should deny that he has, at present, the option of accepting and writing. 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. 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 2.5 if Procrastinate is irrepressible.

So I don't think that cases like *Professor Procrastinate* pose a real worry for maximalism.⁵ But another worry concerning maximalism is that it gives rise to Ross's paradox (1941). To illustrate, imagine that I have promised to mail an important letter for you and so am obligated to do so. Now, mailing the letter entails either mailing or burning it. So maximalism implies that I'm obligated to either mail or burn the letter. Or, at least, it does if we plausibly assume that I'm obligated to perform an option if and only if it is my only permissible option. For, in that case, we can derive from maximalism: S is obligated to perform a non-maximal option if and only if every permissible maximal option entails S's performing that non-maximal option. And every permissible maximal option entails my mailing the letter, which in turn entails my either mailing or burning the letter. Thus, on maximalism, I'm obligated to either mail or burn the letter. Now, some claim that this is false. But they are, I believe, mistaken. Admittedly, it can, in certain contexts, sound infelicitous to claim that I'm obligated to either mail or burn the letter. But we can easily explain this without supposing that the reason for its sounding infelicitous is that it is false. The explanation for its sounding infelicitous is rather that it violates Grice's maxim of quantity, which holds that claiming that I'm obligated to either mail or burn the letter when we know that I'm obligated to mail the letter is inappropriately uninformative.⁶

⁵ There are other objections to maximalism concerning its implication in cases (see Bykvist 2002 and Gustafsson 2014), but, as I argue elsewhere (2015a), they too rely on an implausible view of what it is to have an option now of doing something later on—one according to which I needn't have the present opportunity and ability to influence what I do later on to have now the option of doing that something later on.

⁶ See, for instance, Castañeda 1981 and Wedgwood 2006. But see Cariani 2013 for a criticism of this response to Ross's paradox. Unfortunately, I don't have space here to address Cariani's criticism.

The only other reason that I can see for thinking that “I have an obligation to either mail or burn the letter” must be false (other than it sounds infelicitous) is that when this claim is combined with certain other assumptions, it implies that I have a reason to burn the letter, which is clearly false. For the claim “I have an obligation to either mail or burn the letter” implies that I have a reason to burn the letter if we assume both that (A1) I can fulfill my obligation to either mail or burn the letter by burning the letter and that (A2) if my φ -ing would fulfill an obligation, then that fact itself provides me with a reason to φ . Of course, this means that the extent to which Ross’s paradox should trouble us is in strict proportion to how compelled we should feel to endorse both A1 and A2. And I don’t feel compelled to endorse either, let alone both. I’m happy to say that the only way to fulfill an obligation to perform either φ or ψ in cases where only one of these two is permissible is to perform the permissible one. And I’m also happy to say that the mere fact that my φ -ing would fulfill an obligation is not itself a reason for me to φ —that the only reasons for me to φ are those that account for my being obligated to φ , not the mere fact that my φ -ing is obligatory. So I don’t think that we should reject maximalism on account of its giving rise to Ross’s “paradox.”

Another worry about maximalism is that it entails that we must reject either the validity or the first premise of the following argument, where Repressible Procrastinate is one who would write the review if he were now to respond appropriately to his reasons, accepting the invitation while resolving to get to work on the review as soon as the book arrives.

- (2.15) Repressible Procrastinate is obligated to decline the invitation if he’s going to procrastinate indefinitely. [Intuition]
- (2.16) Repressible Procrastinate is going to procrastinate indefinitely. [From the stipulations of the case]
- (2.17) Therefore, Repressible Procrastinate is obligated to decline the invitation.
(From 2.15 & 2.16)

Maximalism implies that Repressible Procrastinate is permitted (indeed, obligated) to accept the invitation, which contradicts 2.17. So the maximalist must either deny 2.15 or reject the inference from 2.15 and 2.16 to 2.17. Denying 2.16 is not an option, because 2.16 just follows

from the stipulations of the case, which include the stipulation that even if he were to accept the invitation, he would not in fact ever get around to writing the review. Thus, it follows from the stipulations of the case that he is not going to respond appropriately to his reasons and, so, is not going to resolve now to write the review as soon as the book arrives.

Now, it's difficult to deny 2.15, as it seems quite plausible. But I don't think that the maximalist needs to deny 2.15, for she can plausibly deny the validity of the argument. That the argument is invalid can be seen by considering the following structurally identical argument.

(2.18) Olga is obligated to kill Yaakov gently if she's going to kill him. (Intuition)

(2.19) Olga is going to kill Yaakov. (Stipulation)

(2.20) Therefore, Olga is obligated to kill Yaakov gently. (From 2.18 & 2.19)

There are clearly possible worlds in which both 2.18 and 2.19 are true but 2.20 is false, such as the one in which the following are all true: (2.21) Olga is an unfailing hired assassin who's been contracted to kill Yaakov, (2.22) it would be better for Olga to kill him gently as opposed to non-gently, but (2.23) Olga ought not to kill Yaakov, gently or otherwise. Thus, we should think that neither the inference from 2.18 and 2.19 to 2.20 nor the inference from 2.15 and 2.16 to 2.17 is an instance of *modus ponens* or any other valid argument form. We should, then, deny that the conditional expressed by 2.15 is a material conditional with "Repressible Procrastinate is obligated to decline the invitation" as its consequent and "he's going to procrastinate indefinitely" as its antecedent. Instead, we should, perhaps, think that the obligation specified in 2.15 takes wide scope over the entire conditional such that it is saying that Repressible Procrastinate is obligated either to refrain from procrastinating indefinitely or to decline the invitation (Broome 2004). Of course, this is just one possibility on how to interpret the conditional expressed by 2.15.⁷ The maximalist needn't say which possibility is correct, but only that the correct one is such that the inference from 2.15 and 2.16 to 2.17 is invalid. And saying this seems all too plausible.

So far, we've seen no good reason to reject maximalism. There could be other objections that I haven't yet canvassed, but, even so, I'll turn now to make my case against

⁷ Another possibility is that deontic conditionals are a special, primitive type of conditional (see, e.g., von Wright 1956).

maximalism's alternatives. My intent is to make sufficiently strong cases against each of maximalism's alternatives that we'll be more confident that maximalism is true than that one of its alternatives is true.

3. Minimalism

This brings us to the second possibility, *Poss2*, which holds that the permissibility of an option that entails another option is always less fundamental than the permissibility of the option that it entails. I call this view *minimalism*. Like maximalism, it holds that we must evaluate different sorts of options differently. But rather than distinguishing between maximal and non-maximal options (as maximalism does), it distinguishes between minimal and non-minimal options. An option φ is a *minimal option* if and only if there is no option ψ such that φ -ing entails ψ -ing but not vice versa. Otherwise, it's a *non-minimal option*. Thus, typing the word 'the' won't be a minimal option if typing is an option. And typing won't be a minimal option if using a machine is an option. And using a machine won't be a minimal option if doing something is an option. But doing something is a minimal option, for there is no option that it entails but not vice versa. Indeed, the only options that doing something entails are ones that are entailed by it—e.g., performing something. So, even though doing something entails performing something, doing something is still a minimal option given that performing something also entails doing something. Thus, minimal options will be options that are maximally general, such as the option of doing something and the option of performing something.

When fully spelled out, minimalism says the following.

Minimalism: (Min1) For any non-minimal option β , S 's β -ing is permissible if and only if there exists an option φ such that S 's φ -ing is permissible and S 's β -ing entails S 's φ -ing but not vice versa, and when S 's β -ing is permissible, this is in virtue of the fact that there exists an option φ such that S 's φ -ing is permissible and S 's β -ing entails S 's φ -ing but not vice versa. And, (Min2) for any minimal option α , S 's α -ing is permissible if and only if S 's α -ing has feature F , and when S 's α -ing is permissible, this is in virtue of the fact that S 's α -ing has feature F .

According to minimalism, I'm permitted to perform an act so long as it entails performing a type of act that I'm permitted to perform. Thus, I'm permitted to make a bomb so long as I'm permitted to make something. And I'm permitted to feed my daughter poisoned food so long as I'm permitted to feed her some kind of food. I'm even permitted to commit mass genocide so long as I'm permitted to do something. For, on this view, I'm permitted to perform any of my non-minimal options so long as I'm permitted to perform some minimal option—that is, so long as I'm permitted to do something. After all, in each case, my performing that non-minimal option entails my doing something.

This view is, I believe, false. Besides having the counterintuitive verdicts described above, it violates the following plausible principle.

Coarseness Principle: The statement "S is permitted to φ " can be true even though there are many instances of S's φ -ing that are impermissible.

For instance, "I'm permitted to feed my daughter" is true even though there are many instances of feeding her that are impermissible: feeding her dirt, feeding her poisoned food, feeding her shards of glass, etc. But minimalism denies this. On minimalism, if I'm permitted to feed my daughter, because, say, I'm permitted to feed her a nutritious diet, then I'm permitted to perform anything that entails feeding my daughter, including feeding her shards of glass. Thus, I take minimalism to be a nonstarter. I considered it only to exhaust the logical possibilities.

4. Omnism

The next possibility, Poss3, holds that the permissibility of an option that entails another is never either more or less fundamental than the permissibility of the option that it entails. I call this view *omnism*, because 'omnis' means 'all' in Latin and this view holds that the permissibility of all options, whether they entail or are entailed by others, are equally fundamental. On this view, there is no need to distinguish between different types of options (such as, between minimal and non-minimal options or between maximal and non-maximal options). For, on omnism, all options are to be evaluated in the same way. When fully spelled out, the view says the following.

Omnism: For 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 .

Unlike maximalism and minimalism, omnism is an immensely popular view. Virtually all theories of morality and rationality that have been offered to date are versions of omnism.⁸ Substitute 'maximizes utility' for 'has feature F ', and we get act-utilitarianism. Substitute 'accords with Kant's categorical imperative' for 'has feature F ', and we get Kantianism. And substitute 'maximizes S 's utility' for 'has feature F ', and we get rational egoism. But despite its popularity, I believe that we should reject omnism. As I'll show, several versions of omnism offer an implausible account of what makes maximal options permissible, and all others problematically generate conflicting obligations—instances in which a subject is both obligated to φ and obligated to refrain from ψ -ing even though it is logically impossible for her both to φ and to refrain from ψ -ing given that her φ -ing logically necessitates her ψ -ing.

To see that some versions generate conflicting obligations, consider act-utilitarianism's implications in the following case.

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 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

⁸ As far as I know, the only exceptions are the theories defended in Feldman 1986, Goldman 1978, Portmore 2011, and Zimmerman 1996, which are all versions of maximalism.

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.

Act-utilitarianism implies that Doc is obligated to inject Pat with a proper dose of D, for this would produce more utility than any alternative would.⁹ But, interestingly, act-utilitarianism also implies that Doc is obligated to refrain from injecting Pat with D, because injecting Pat with D would produce less utility than the alternative of refraining from injecting him with D would. It's stipulated, after all, that if Doc were to inject Pat with D, she would inject Pat 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 the alternative of refraining from injecting him with D. Thus, act-utilitarianism implies both that Doc is obligated to inject Pat with a proper dose of D and that Doc is obligated to refrain from injecting him with D, and yet it is logically impossible for Doc both to inject him with a proper dose of D and to refrain from injecting him with D, for injecting him with a proper dose of D logically necessitates injecting him with D. And, so, we see that act-utilitarianism generates conflicting obligations.

Such conflicting obligations are problematic, because the sorts of obligations at issue are meant to be action-guiding. That is, the sense of 'obligation' at issue is the one where, as a matter of conceptual necessity, believing that one has an obligation (in this sense) to φ rationally requires one to form the intention to φ . Thus, as I understand them, theories such as omnism, maximalism, and minimalism just are theories about obligations for which the following principle holds.

Enkratic Principle: Rationality requires that a subject be such that, if she believes that she's obligated to φ , she intends to φ (Broome 2013).

⁹ Act-utilitarianism holds that, for any option φ , S 's φ -ing is permissible if and only if there is no alternative option ψ such that S 's ψ -ing would produce more utility than S 's φ -ing would. 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-utilitarianism, an act is obligatory if it would produce more utility than any alternative would.

Given that this is the sense of ‘obligation’ at issue, any version of omnism that generates conflicting obligations is committed to the denial of the plausible idea that rationality demands “that one’s intentions, taken together with one’s beliefs, fit together into a consistent model of one’s future” (Bratman 2009, p. 29). More specifically, we’ll have to deny the following.

Belief-Intention Consistency Principle: 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 .¹⁰

And here’s my argument for why we must deny conflicting obligations if we are to accept the above two principles.¹¹

- (4.1) S is both obligated to φ and obligated to refrain from ψ -ing even though it is logically impossible for her both to φ and to refrain from ψ -ing. (Assumption for *reductio*)
- (4.2) Rationality requires that S be such that, if she believes that she’s obligated to φ , she intends to φ . (The Enkratic Principle)
- (4.3) Thus, rationality requires that, if S believes what’s true about what she’s obligated to do, she both intends to φ and intends to refrain from ψ -ing. (From 4.1 and 4.2)
- (4.4) If it is logically impossible for S both to φ and to refrain from ψ -ing, then she will not both φ and refrain from ψ -ing. (Analytic)
- (4.5) Thus, if S believes what’s true about what she will do, she believes that she will not both φ and refrain from ψ -ing. (From 4.1 and 4.4)
- (4.6) Thus, if S is as rationality requires and believes what’s true about what she will and is obligated to do, she both intends to φ and intends to refrain from

¹⁰ The variables φ and ψ range over both acts (e.g., x) and omissions (e.g., refrain from x -ing). This is the contrapositive of what Jacob Ross calls the *Strong Belief Requirement*—see his 2009, p. 246. And note that this is not the metaphysical view that intending to φ involves believing that one will φ . Rather, it is the normative view that intending to φ rationally requires one to believe that one will φ . For a defense, see Bratman 1984 and 1987.

¹¹ This argument comes from Kieseewetter 2015 (pp. 929–934). It is along similar lines to, though much better formulated than, the one that I gave in Portmore 2011 (pp. 181–183). It is, I believe, what I was grasping at then.

ψ -ing and believes that she will not both φ and refrain from ψ -ing. (From 4.3 and 4.5)

- (4.7) Rationality requires that S be such that, if she believes that she will not both φ and refrain from ψ -ing, then she does not both intend to φ and intend to refrain from ψ -ing. (From the Belief-Intention Consistency Principle)
- (4.8) Thus, if S is as rationality requires and believes what's true about what she will and is obligated to do, she will not be as rationality requires. (From 4.6 and 4.7)
- (4.9) It is not the case that, if S is as rationality requires and believes what's true about what she will and is obligated to do, she will not be as rationality requires. (Intuition)
- (4.10) Therefore, 4.1 is false; there are no conflicting obligations. (From 4.1, 4.8, and 4.9)

So, given that omnism is about obligations in the sense for which 4.2 is a conceptual necessity, the versions of omnism that generate conflicting obligations (e.g., act-utilitarianism) must reject the Belief-Intention Consistency Principle. But we should not reject this plausible principle; we should instead reject any version of omnism that generates conflicting obligations.¹²

It's not just act-utilitarianism that generates conflicting obligations in *The Injection*. Any version of omnism that holds that an option is F if and only if it has the feature of being unsurpassed in terms of justice, fidelity, non-maleficence, respect for autonomy, overall utility production, production of utility for the agent, etc. will generate conflicting obligations in this sort of case. And, to save words, I'll hereafter call this feature 'G'. Thus, my claim is that any version of omnism that holds that an option is F if and only if it is G will generate conflicting obligations. For such versions of omnism will entail both that (OB1) Doc is obligated to inject Pat with a proper dose of D and that (OB2) Doc is obligated to

¹² Although what I've said rules out the possibility of conflicting moral obligations in the action-guiding sense of 'moral obligation' for which the Enkratic Principle applies, it does not rule out the possibility of moral dilemmas in which an agent will do something for which she will be morally blameworthy regardless of what she does, as where she has separately and negligently made two promises that she cannot both keep.

refrain from injecting Pat with D. This is because, as I'll now explain, injecting Pat with a proper dose of D is Doc's only option that is G.

Admittedly, given the above description of *The Injection*, it could be that Doc's injecting Pat with a proper dose of D isn't even G—let alone her only option that is G. For, as far as the above description goes, it could be that Doc has the option of killing Pat with an improper dose of D and then using his stem cells to synthesize a drug that would save thousands of people of Pat's race. And if this were an option, injecting him with a proper dose of D would not be G, since injecting him with an improper dose of D and then using his stem cells to save thousands would produce more overall utility.¹³ But, for my purposes, I want a case in which Doc has no such option—that is, I want a case in which Doc has no better option than injecting Pat with a proper dose of D. So I hereby add to the above description the stipulation that there is no such option and that, therefore, Doc's injecting Pat with a proper dose of D is G and, indeed, her only option that is G.

Adding this stipulation ensures that any version of omnism that holds that an option is F if and only if it is G will entail both that (OB1) Doc is obligated to inject Pat with a proper dose of D and that (OB2) Doc is obligated to refrain from injecting Pat with D. It will entail OB1 given that Doc's injecting Pat with a proper dose of D is her only option that is G and, thus, her only permissible option.¹⁴ Moreover, it will entail OB2, because injecting Pat with D is not G and, therefore, is not permissible. And if it's impermissible for Doc to inject Pat with D, then she is obligated to refrain from injecting him with D.¹⁵ Thus, any version of omnism that holds that an option is F if and only if it is G will generate both OB1 and OB2 in *The Injection*. This includes not only consequentialist versions of omnism like act-utilitarianism but also non-consequentialist versions of omnism like Rossian pluralism. So it seems that a great many versions of omnism will generate conflicting obligations.

The only versions of omnism that won't generate conflicting obligations are those that deny that an option is F if and only if it G. And there are only two ways that the omnist might do this. The first is to specify her F so that it implies that OB1 is false even though

¹³ Now, this may seem incompatible with the above stipulation that injecting Pat with an improper dose of D would produce less overall utility than injecting him with a proper dose of D. But it isn't, for it could be that, given Doc's hatred towards people of Pat's race, she would not use Pat's stem cells to save thousands of people of Pat's race even if she were to kill him with an improper dose of D.

¹⁴ I'm assuming that S is obligated to ϕ if and only if ϕ is S's only permissible option.

¹⁵ I'm assuming that S is prohibited from ϕ -ing if and only if S is obligated to refrain from ϕ -ing.

Doc's injecting Pat with a proper dose of D is her only option that is G. The second is to specify her F so that although OB1 is true, OB2 is false—false even though injecting Pat with D is not G. I'll argue that both ways are problematic, starting with the first.

To specify her F so that it implies that OB1 is false, the omnist must substitute for 'has feature F' something that doesn't correlate with G—something like 'does not entail the commission of murder'. The result is what I'll call *no-murder omnism*. On no-murder omnism, OB1 is false despite the fact that Doc's injecting Pat with a proper dose of D is her only option that is G. It's false, because Doc is permitted to refrain from injecting Pat with D given that so refraining doesn't entail the commission of murder. The problem, though, is that it's implausible to deny that Doc has an obligation to inject Pat with a proper dose of D. And that's because Doc's injecting him with a proper dose of D is her only option that is G and, thus, her best option. Indeed, it's her best option by far. It's certainly her best option in terms of fidelity given that doing anything else would entail violating her Hippocratic Oath. It's also her best option in terms of both overall utility and utility for herself, for it's the only option that results in everyone's living happily ever after. And it's her best option in terms of respect for autonomy given that Pat has, we'll assume, autonomously requested to receive a proper dose of D. Lastly, injecting him with a proper dose of D seems to be unsurpassed in terms of any other possibly relevant consideration. Thus, injecting Pat with a proper dose of D is Doc's only option that is G and, thus, her best option. And it's implausible to deny that Doc has an obligation to perform her best option in circumstances like these where there are multiple and significant factors at stake.¹⁶ And this means that the first way of avoiding conflicting obligations is implausible. For it involves specifying the omnist's F so that it implausibly entails that OB1 is false.

This brings us to the second way for omnism to avoid conflicting obligations in *The Injection*, and that is to specify its F such that OB1 is true but OB2 is false. To get the one obligation without the other, the omnist must specify her F such that some options (e.g., injecting Pat with a proper dose of D) are F if and only if they are G but that other options (e.g., injecting Pat with D) are F even though they are not G. And to do that she must jury-

¹⁶ I would allow that there could be cases in which there is no obligation to φ even though φ is the only option that is G, but these would all have to be cases in which φ is the only option that is G because φ is just slightly better in terms of utility. But, in *The Injection*, injecting Pat with a proper dose of D is significantly better in terms of utility. Moreover, it is better in terms of more than just utility—it is, for instance, also better in terms of fidelity and respect for autonomy.

rig F. That is, she must substitute something like 'is entailed by a maximal option that is G*' for 'has feature F' in the above formulation of omnism, where G* is any feature that necessarily correlates with G. I'll call the resulting theory *maximalist omnism*, because it is extensionally equivalent to the version of maximalism that substitutes 'has feature G*' for 'has feature F' in the above formulation of maximalism. Now, on maximalist omnism, Doc's injecting Pat with a proper dose of D is obligatory, because it is entailed by all maximal options that are G*. After all, the only maximal options that are G* are those in which Doc injects Pat with a proper dose of D. And, naturally, all these entail Doc's injecting Pat with a proper dose of D. Thus, maximalist omnism, unlike no-murder omnism, plausibly endorses OB1. But it denies OB2, thereby avoiding conflicting obligations. For, on maximalist omnism, Doc is permitted to inject Pat with D. What's more, Doc is obligated to inject Pat with D. For, again, all the maximal options that are G* entail injecting Pat with a proper dose of D and, thus, also entail injecting him with D. Thus, even though injecting Pat with D is not G it is, on maximalist omnism, obligatory. And so maximalist omnism entails OB1 but not OB2.

The problem with this way of avoiding conflicting obligations, though, 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 G*. Rather than simply claiming that performing such a maximal option is permissible because it is G*, the omnist must claim something like one of the following: (E1) it is permissible because it is entailed by a maximal option that is G*; (E2) it is permissible because it is entailed by a maximal option that is entailed by a maximal option that is G*; (E3) it is part of the ideal future in which she performs only maximal options that are G*; (E4) it is part of the ideal future in which all agents perform only maximal options that are G*; etc.

To illustrate the problem, suppose that 'has feature G*' stands for something like 'is in accordance with Kant's categorical imperative' or 'is a what a virtuous person would characteristically do in the circumstances' — assuming, of course, that these necessarily correlate with G, and insofar as they are plausible it seems that they must. Now, the omnist who is adopting this approach to avoiding conflicting obligations cannot simply substitute 'has feature G*' for 'has feature F' in the formulation of omnism, for, given that G* necessarily correlates with G, this would commit her to both OB1 and OB2. Instead, she

must specify her F so that maximal options will be permissible if and only if they are G*, but that non-maximal options will be permissible provided that they are entailed by a maximal option that is G* and, thus, regardless of whether they are themselves G*. Thus, she has to jury-rig her F so that the resulting verdicts are extensionally equivalent to the version of maximalism that substitutes G* for maximalism's F. But although this allows the omnist to get the same verdicts, it commits her to a different account of what makes maximal options permissible. Whereas the maximalist will hold that a permissible maximal option is permissible in virtue of its having feature G*, the omnist must hold that a permissible maximal option is permissible in virtue of something like its being entailed by a maximal option that is G*. The problem, then, is that this seems like an implausible explanation for why it's permissible to perform a maximal option that is G*. Whether a maximal option is G* seems relevant. But whether a maximal option is entailed by a maximal option that is G* 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 G*. Likewise, whether a maximal option is part of the ideal future in which the agent performs only maximal options that are G* seems no more relevant than whether it is part of the ideal future in which all agents perform only maximal options that are G*. Thus, the second way of avoiding conflicting obligations forces the omnist to adopt an implausible view about what makes maximal options that are G* permissible.

So I've argue that, in order to avoid conflicting obligations, the omnist must either (1) specify her F so that it implies that OB1 is false or (2) specify her F so that although OB1 is true, OB2 is false. And I've argued that both are implausible. The first is implausible because it's implausible to deny OB1, and the second is implausible because it's implausible to deny that what makes a maximal option that is G* permissible is that it is G*. Thus, we should, I believe, reject omnism.

5. Hybridism

This brings us to the last of the four possibilities: Poss4. On this view, there is a lack of uniformity, as there would be, for instance, if sometimes the permissibility of an option that entails another is more fundamental than the permissibility of the option that it entails, but other times the permissibility of an option that entails another is less fundamental than the

permissibility of the option that it entails. For obvious reasons, I'll call this view *hybridism*. Now, there are many different possible versions of hybridism, but I'll consider just one. For it seems rather straightforward how the problems for it—problems involving both arbitrariness and inheriting the vices of at least one of its two parents—easily generalize to other possible versions of hybridism. The version of hybridism that I'll consider is the one that seems most natural. It holds that whether we take the permissibility of a given option to be a function of the options that it entails, a function of the options that entail it, or a function of neither depends on the given option's level of entailment, where an option's level of entailment is a function of both which other options it entails and which other options entail it, such that maximal options (which entail all other options but are entailed by no other option) have the highest level of entailment and minimal options (which entail no other option but are entailed by every other option) have the lowest level of entailment. When fully spelled out, the view is as follows.

Hybridism: (Hyb1) For any option z whose level of entailment is \mathcal{E} , S 's z -ing is permissible if and only if S 's z -ing has feature F , and when S 's z -ing is permissible, this is in virtue of the fact that S 's z -ing has feature F . (Hyb2) For any option x whose level of entailment is greater than \mathcal{E} , S 's x -ing is permissible if and only if there is an option z such that (a) it is F , (b) S 's x -ing entails S 's z -ing, and (c) its level of entailment is \mathcal{E} , and when S 's x -ing is permissible, this is in virtue of the fact that there is such a z . And (Hyb3) for any option y whose level of entailment is less than \mathcal{E} , S 's y -ing is permissible if and only if there is an option z such that (a) it is F , (b) S 's z -ing entails S 's y -ing, and (c) its level of entailment is \mathcal{E} , and when S 's y -ing is permissible, this is in virtue of the fact that there is such a z .

Hybridism seems to be the least promising of the four possibilities. For one, it seems that unless we set \mathcal{E} at one of the two extremes (that is, at either the maximal or minimal level of entailment), it will be arbitrary, as there will be no justification for setting it at that point and not at some point slightly higher or lower than it. And if we do set \mathcal{E} at one of the two extremes, it will cease to be distinct, for it will collapse into either maximalism or minimalism. If, on the one hand, we set \mathcal{E} at the maximal level of entailment, which is the

level that only maximal options reach, then there will be no option x whose level of entailment is greater than \mathcal{E} . And, in that case, Hyb2 drops out, and hybridism collapses into maximalism. If, on the one hand, we set \mathcal{E} at the minimal level of entailment, which is the level that all but minimal options exceed, then there will be no option y whose level of entailment is less than \mathcal{E} . And, in that case, Hyb3 drops out, and hybridism collapses into minimalism. So either hybridism is going to be arbitrary or it's going collapse into either maximalism or minimalism.

And this is not hybridism only problem. In holding Hyb2, hybridism becomes sufficiently like minimalism to be subject to the same sorts of counterexamples that it's subject to. To illustrate, let's again consider *The Injection*. And let's suppose that the level of entailment of Doc's injecting Pat with the proper dose of D is precisely at \mathcal{E} . In setting \mathcal{E} at this point, the hybridist can avoid implying both that (OB1) Doc is obligated to inject Pat with the proper dose of D and that (OB2) Doc is obligated to refrain from injecting Pat with D. Since injecting Pat with D is at a lower level of entailment than injecting Pat with the proper dose of D, we do not assess it in terms of F. Instead, we assess it in terms of whether it is entailed by a permissible option that is at level \mathcal{E} . And, indeed, it is entailed by injecting Pat with the proper dose of D, which is an option at level \mathcal{E} . Thus, hybridism implies that Doc is obligated to inject Pat with D. So there's no conflict of obligations here. But, given that hybridism accepts Hyb2, we have to hold that any option y that's at a greater level of entailment than \mathcal{E} (such as, the option of injecting Pat with the proper dose of D and then slitting his throat afterwards) is permissible if there is an option z such that (a) it is F, (b) it is at level \mathcal{E} , and (c) it is entailed by the performance of y . And there is such an option z —namely, the option of injecting Pat with the proper dose of D. Thus, like minimalism, hybridism implies such absurd things as that Doc is permitted to inject Pat with the proper dose of D and then slit his throat afterwards. Given this, hybridism seems as much of a non-starter as minimalism.

These problems will generalize to any version of hybridism. Even if we use something besides the level entailment to determine which options are to be assessed in terms of F, which options are to be assessed in terms of options that entail it, and which options are to be assessed in terms of options it entails, we're going to run into these

problems of being arbitrary and/or inheriting minimalism's counterintuitive implications. So I believe we should reject hybridism.

6. Maximalism's Explanatory Component

Maximalism (i.e., Poss1) seems to be the clear winner, for it seems to be the only possibility without unsurmountable difficulties. Some, however, may worry that I've overlooked an important difficulty with maximalism. Specifically, some may worry that maximalism gives the wrong explanation for why permissible non-maximal options are permissible. According to maximalism, if a subject is permitted to perform a non-maximal option ν , this is in virtue of the fact that her ν -ing is entailed by a maximal option that has the right-making feature F. But it may seem that the explanation for why she is permitted to ν lies simply with the fact that her ν -ing is itself F 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. But, surely, 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 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 brush my teeth without following up by drinking a glass 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 that 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. 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 his patient, Prisha, at t_2 . In fact, there are exactly two drugs available to him: A and B. If he gives Prisha both A and B at t_2 , using his right hand to inject her with A and his left hand to inject her with B, she'll be cured of her painful, but non-fatal, medical condition. If he gives her just one of A and B at t_2 , she'll die immediately. And if he gives her neither A nor B at t_2 , she'll be rendered incurable, having, then, to live the rest of her life with this painful medical condition. And although this would be quite bad, it would not be as bad as her dying immediately. Now, as a matter of fact, Dr. Singh is going to form at t_1 the intention to kill her by giving her just one of A and B at t_2 . Moreover, he's going to follow through with this intention. For, as it turns out, he is Prisha's jilted lover. Given that he's going to form, and then follow through, with the intention to give her just one of A and B at t_2 , the following two subjunctive conditionals are true: (SC1) if he were to give her A at t_2 , he would not give her B at t_2 and (SC2) if he were to give her B at t_2 , he would not give her A at t_2 . Nevertheless, if he were to respond appropriately to the decisive reason he has for curing her and thereby come instead to form at t_1 the intention to cure her by giving her both A and B at t_2 , he would follow through with this intention and Prisha would live happily

ever after. Lastly, assume that Dr. Singh will lose his medical license if he fails to cure her and will go to prison if he kills her.

It seems that Dr. Singh is both permitted to give Prisha A at t_2 and permitted to give her B at t_2 . But the explanation for why this is so has nothing to do with facts about each of these acts. For instance, it's not that performing each of these act would have good consequences. In fact, given the truth of SC1, injecting her with A at t_2 would have terrible consequences. And, given the truth of SC2, injecting her with B at t_2 would also have terrible consequences. Instead, the explanation for why Dr. Singh is both permitted to give her A at t_2 and permitted to give her B at t_2 is that he's permitted to do something that entails doing each of these. That is, he is permitted to give Prisha both A and B at t_2 , and this logically necessitates both giving her A at t_2 and giving her B at t_2 . Taking this further, it would seem that the explanation for why she is permitted give Prisha both A and B at t_2 is that he is permitted to perform some option that entails doing this, such as injecting her with these 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 Prisha with A, not because doing so would itself have optimal consequences (it wouldn't), but because it is entailed by some maximal option that would have optimal consequences.

7. Conclusion

I've argued that we should accept maximalism. It encapsulates 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, as I've showed, none of the objections that have been offered against maximalism are persuasive. And, in this respect, maximalism is superior to its three alternatives (viz., omnism, hybridism, and minimalism). All of which have serious problems.

If maximalism is true, then nearly all the theories of morality and rationality that have so far been proffered are false, for nearly all of them are versions of omnism. They are versions of omnism in that they don't even distinguish between maximal and non-maximal options. And in failing to distinguish between these two types of options, they presume that we can assess the deontic statuses of non-maximal options without appealing to the deontic statuses of the maximal options that entail them. That is, they assume that we can evaluate both maximal and non-maximal options in terms of whether or not they have some right-making feature F. As a consequence, they will either generate conflicting obligations or offer an implausible explanation for what makes permissible maximal option permissible.

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 actions, and, for all that I've said, they may have correctly done so. They just make the mistake in 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.¹⁷

Arizona State University
dwpportmore@gmail.com

References

- Bratman, M. E. (2009). "Intention, Belief, Practical, Theoretical." In S. Robertson (ed.), *Spheres of Reason*, pp. 29–62. Oxford: Oxford University Press.
- . (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- . (1984). "Two Faces of Intention." *Philosophical Review* 93: 375-405.
- Broome, J. (2013). *Rationality through Reasoning*. Oxford: Wiley-Blackwell.

¹⁷ For helpful comments and questions, I thank Seth Lazar, Felix Pinkert, Thomas Mautner, and audiences at Australian National University and Charles Sturt University, Wagga Wagga. I especially thank Wylie Breckenridge, Daniel Cohen, and N. 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.

- . (2004). "Reasons." In R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, pp. 28–55. Oxford: Oxford University Press.
- 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. (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. and R. Pargetter (1986). "Oughts, Actions, and Actualism." *Philosophical Review* 95: 233–255.
- Kiesewetter, B. (2015). "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125: 921–946.
- Portmore, D. W. (2015a). "Maximalism and Rational Control." Working manuscript that available at <http://bit.ly/1NmZ57E>.
- . (2015b). "Maximalism and Supererogation." Working manuscript.
- . (2013). "Perform Your Best Option." *The Journal of Philosophy* 110: 436–459.
- . (2011). *Commonsense Consequentialism: Wherein morality meets rationality*. New York: Oxford University Press.
- Ross, J. (2009). "How to Be a Cognitivist about Practical Reason." *Oxford Studies in Metaethics* 4: 243–282.
- 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.
- von Wright, G. H. (1956). "A Note on Deontic Logic and Derived Obligation." *Mind* 65: 507–509.
- 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.