

What are our options?

DOUGLAS W. PORTMORE

ABSTRACT: We ought to perform our best option—that is, the option that we have most reason, all things considered, to perform. This is perhaps the most fundamental and least controversial of all normative principles concerning action. Yet, it is not, I believe, well understood. For even setting aside questions about what our reasons are and about how best to formulate the principle, there is a question about how we should construe our options.¹ This question is of the utmost importance, for which option will count as best depends on how broadly or narrowly our options are to be construed. In this paper, I argue that an agent's options at a time, *t*, are all and only those actions (or sets of actions) that are scrupulously securable by her at *t*.

WE OUGHT TO do our best. That is, we ought to perform our best option. By 'best option', I don't mean the option that would produce the best outcome. The idea that we ought to do what would produce the best outcome is quite controversial. Instead, I mean something much less controversial: we ought to perform the option that we have most reason (all things considered) to perform.² Of course, there could be more than one option that's tied for first-place. So, to be a bit more careful, I should say that we ought to perform one of the options that we have optimal reason to perform.

¹ I address the issue of how best to formulate the principle in a companion paper entitled "Perform Your Best Option" (2011b). Both this paper and its companion are descendants of a now defunct paper entitled "Doing Our Best," which is cited in Portmore 2011a. In a separate paper, I address the issue of what our reasons are—see Portmore 2011c.

² Throughout, I'll be concerned, not with what agents morally ought to do or with what agents prudentially ought to do, but with what agents ought to do, all things considered. Moreover, I'll be concerned with only what agents *objectively* ought (or are permitted or are obligated) to do—that is, with what they ought (or are permitted or are obligated) to do given what the relevant reason-constituting facts about their choice situation happen to be, and so irrespective both of what they take those facts to be and of what their evidence suggests that those facts might be. And when I talk about reasons, I'll be talking about *objective* reasons—these are facts that count in favor of an agent's performing a given action irrespective of her beliefs or evidence. For instance, the fact that there will be dancing at the party constitutes a reason for Ronnie, who enjoys dancing, to attend the party regardless of whether or not he believes, or has any reason to believe, that there will be dancing at the party (Schroeder 2007). For more on these distinctions, see Portmore 2011a, pp. 12–23.

Now, the word ‘optimal’ is a comparative term. To say that an agent has optimal reason to perform a given option is to say that there is no alternative option that she has more reason to perform. But what are the agent’s options? There seem to be a number of ways that we might construe an agent’s options: those actions that it is logically possible for her to perform, those actions that it is physically possible for her to perform, those actions that it is personally possible for her to perform, etc.³ For now, though, let me sidestep the issue by stipulating that I’ll use the term ‘ x options’ to stand for whatever the relevant options are. That is, we are to substitute for ‘ x ’ (perhaps, ‘personally possible’) whatever would make the following maximally plausible:

PYB x O A subject, S , is permitted at t to perform an action, ϕ , at t' if and only if (and because) ϕ -ing at t' is an x option for S at t and there is no other action that is an x option for S at t that S has more reason to perform ($t < t'$).

I call this “PYB x O” (pronounced: PAHY-biks-oh), because it implies that you are obligated to *perform your best x option*—that is, the action that, of all your x options, is the one that you have most reason to perform.⁴

According to PYB x O, what an agent can be obligated to do is constrained by what her x options are. If her ϕ -ing isn’t an x option, then she can’t be obligated to ϕ . But what was once an x option for an agent may no longer be one. To illustrate, imagine that, last week, Jane had the x option of enlisting in either the Army or the Navy. But now that she has enlisted in the Army, she no longer has the x option of enlisting in the Navy, for she can’t enlist in the Navy if she’s already enlisted in the Army. (Assume, for now, that ϕ is an x option for S only if S can ϕ .) And if her x options can change over time, then so too can her obligations—at least, if PYB x O is true. Suppose, for instance, that Jane had promised her father a month ago that she would enlist in the Navy on her birthday, which is today. A month ago, then, she had an obligation to enlist in the Navy on her birthday. But, as of yesterday, when she enlisted in the Army, she no longer has the x option of enlisting in the Navy. Thus, she no longer has an obligation to do so. Of course, she may have an obligation to apologize to her father for breaking her promise. But she cannot now be obligated to do what isn’t even an x option for her—at least, not if PYB x O is true. This means that, if we’re

³ I borrow the term ‘personally possible’ from Zimmerman 2007, and I define it in section 3 below.

⁴ It implies this, for S is obligated to ϕ if and only if ϕ -ing is S ’s only permissible x option.

careful, we should, as I've done above, make explicit the relevant temporal indices when formulating PYBxO.

If it seems strange to you to talk about permissions and options that are indexed to times, then think of them as I do: as properties that are possessed by the agent at certain times. Thus, you can think of the phrase "S is permitted at t to ϕ at t " as equivalent to the phrase "S has at t the property of being permitted to ϕ at t ." And, likewise, you can think of the phrase " ϕ -ing at t is an option for S at t " as equivalent to the phrase "S has at t the property of having ϕ -ing at t as an option." Call the first sort of property a deontic property and the latter, an option property. Clearly, people's option properties can change over time. And, given PYBxO, so must their deontic properties.

In this paper, I'll address the issue of how we should construe an agent's x options so as to make PYBxO maximally plausible. This will involve three steps. First, in section 1, I'll explain how deontic principles differ from evaluative principles. Second, in section 2, I'll specify three conditions that a deontic principle (such as PYBxO) must meet in order to be plausible. Third, in section 3, I'll consider various alternative proposals for what we might substitute for x and assess the resulting versions of PYBxO according to the conditions specified in section 2. The upshot will be that we should construe an agent's x options at t to be all and only those actions (or sets of actions) that are *scrupulously securable* by her at t . That is, we should substitute 'scrupulously securable' for ' x ' in PYBxO and arrive at PYBSSO (PAHY-bis-soh): the view that you ought to *perform your best scrupulously securable option*. In section 4, I'll argue that PYBSSO is just right in that it construes our options neither too broadly nor too narrowly. Lastly, in section 5, I'll address some potentially problematic cases for PYBSSO.

1. Deontic vs. Evaluative Principles

For each way of construing an agent's options, there is a different version of PYBxO. That is, we get a different version of PYBxO depending on what we substitute for ' x ' in the above schema. For instance, if we construe an agent's options as consisting in all and only those acts that are logically possible, thereby substituting 'logically possible' for ' x ' in the above schema, we get a version of PYBxO that holds that you ought to *perform your best logically possible option* (i.e., 'PYBLPO'—PAHY-blip-oh). So, if we want to determine how to construe an agent's x options so as to make PYBxO maximally plausible, we must look at the plausibility of the resulting versions of PYBxO. If, for instance, we want to determine the plausibility of construing an agent's options so as to include all and only those acts that are logically possible, we must assess the plausibility of

PYBLPO. And, here, it will be helpful to understand how deontic principles (such as PYBLPO) differ from evaluative principles. For this, in turn, will help us to understand what we should be looking for in a deontic principle and thereby allow us to settle on some conditions for their plausibility, conditions that we can then use to assess the various possible versions of PYBxO.

The key to understanding how deontic principles differ from evaluative principles is to understand that there are two ways that we can assess a given action. We can assess not only whether it would be fitting to intend (or to choose) to perform the action, but also whether it would be fitting to desire that the action be performed.⁵ To assess the former, we must employ some deontic principle—that is, a criterion of rightness that allows us to classify acts according to their deontic status: obligatory, impermissible, or optional.⁶ And, to assess the latter, we must employ some evaluative principle—that is, a criterion of goodness that allows us to classify acts according to their evaluative status: good, bad, or neutral (and also: better or worse than some alternative act).⁷

Whereas we can sensibly ask of any act (indeed, of anything at all, including any object, event, or state of affairs) whether it is fit to be desired, we cannot sensibly ask of any act whether it is fit to be intended. To illustrate, consider the

⁵ I take the notion of fittingness to be a primitive notion. The best way to get a grip on it, then, is by considering examples. The fact that intending now to drink some toxin tomorrow will result in one's receiving a million dollars may in some sense count in favor of intending now to drink the toxin tomorrow, but it doesn't make it fitting to intend now to drink the toxin tomorrow. To the contrary, the fact both that drinking the toxin tomorrow will make no difference to whether or not one receives the million dollars and that drinking the toxin tomorrow will only cause one to become ill makes it fitting to intend now *not* to drink the toxin tomorrow. Or, to take another example, the fact that one would receive a million dollars for believing that grass is purple may in some sense count in favor of believing that grass is purple, but it doesn't make it fitting to believe that grass is purple. It is the fact that grass looks green (and not purple) that makes it fitting to believe that grass is green rather than purple. For an excellent discussion of these issues, see Chappell *forthcoming*, from which I borrow the grass example. The toxin example is borrowed from Kavka 1983.

⁶ An obligatory act is one that it would be unfitting not to intend to perform, an impermissible act is one that it would be unfitting to intend to perform, and an optional act is one that it would be neither unfitting to intend not to perform nor unfitting to intend to perform.

⁷ A good act is one that it would be unfitting not to desire to be performed, a bad act is one that it would be unfitting to desire to be performed, and a neutral act is one that it would be neither unfitting to desire to be performed nor unfitting to desire not to be performed. And an act, a_i , is better than an alternative act, a_j , if and only if it would be unfitting not to prefer the performance of a_i to the performance of a_j .

following example, which I'll call *Lake*. Suppose that the only way for Abe to get to his friend and financier, Fred, in time to save his life is to run across the water of the lake that separates them. Given that Fred will otherwise die, it certainly seems fitting for Abe (and for everyone else) to desire that Abe runs on water and saves Fred. What decent person would not want this? But is it also fitting for Abe to intend to run across the lake when he knows that such an intention would be ineffective given his inability to run on water? It seems not. For it seems that only those acts that are under the agent's control can be rationally intended. And, thus, it is only those acts that are under the agent's control that are open to deontic assessment.

We might wonder, though, why deontic principles can be used to assess only those acts that are under the agent's control, whereas evaluative principles can be used to assess anything at all. Furthermore, we might wonder why deontic principles aren't simply superfluous given that evaluative principles can be used to assess anything that deontic principles can assess and more. That is, we might wonder why it is not enough simply to assess acts in terms of their goodness. This question seems all the more pressing given that a ranking of acts in terms of their evaluative statuses will be much more fine-grained than a ranking of acts in terms of their deontic statuses. Although we can surely rank any obligatory act above any optional act and any optional act above any impermissible act in terms of how much reason the agent has to perform them, we can get an even more fine-grained ranking of actions if we appeal to the evaluative statuses of these acts. After all, some optional acts are better than others—for instance, even if giving anything over 10% of one's income to charity is optional, it may be better to give 20% as opposed to 15%. That is, even if both acts would be optional, it may be that the agent has more reason to give 20% than to give 15%. And some impermissible acts are worse than others—for instance, killing oneself is probably worse than merely chopping off one's pinky finger. Both may be impermissible, but there is even greater reason to refrain from killing oneself than there is to refrain from chopping off one's pinky—assuming that there's nothing wrong with either one's life or one's pinky. So, again, we might wonder why deontic principles aren't simply superfluous given that we get a more fine-grained ranking of actions using evaluative principles.⁸

⁸ Indeed, some philosophers argue that we can dispense with deontic principles altogether. For instance, Norcross claims that "once a range of options has been evaluated in terms of goodness, all the morally relevant facts about those options have been discovered." He claims that there is no further fact about which of these options are obligatory, which of them are optional, and which of them are impermissible. See Norcross 2006, p. 44. Thus, as Norcross sees things, deontic principles are superfluous,

To illustrate the worry, imagine that there is a restaurant that serves beverages by the fluid ounce. You can specify whether you want, say, an eight ounce beverage, a nine ounce beverage, or a ten ounce beverage. And suppose that they also categorize beverages into three coarse-grained categories: small (one to ten ounces), medium (eleven to twenty ounces), and large (twenty-one ounces and up). But despite this, you still have to order some specific number of ounces. It seems, then, that the coarse-grained small-medium-large ranking is completely superfluous given the fine-grained one-two-three-and-so-forth ranking. The worry about deontic principles, then, can be expressed as follows: given the fine-grained this-or-that-much-reason-to-perform ranking, why isn't the coarse-grained obligatory-optional-impermissible ranking of actions superfluous in exactly the way that the small-medium-large ranking of beverage sizes is superfluous given the fine-grained one-two-three-and-so-forth ranking of beverage sizes? This calls out for an explanation.

There are actually two things that need explaining: (1) why deontic principles can be used to assess only those acts that are under the agent's control when evaluative principles can be used to assess anything at all and (2) why deontic principles are anything but superfluous given that evaluative principles allow us to give a more fine-grained ranking of actions in terms of how much reason there is to perform them. The explanation in each case lies, I believe, with the fact that deontic principles, unlike evaluative principles, fulfill certain roles beyond the mere classificatory role of putting actions into various categories (such as obligatory, optional, impermissible, better than, worse than, etc.) and thereby ranking them according to these categories. These additional roles are the key to understanding why deontic principles are not superfluous and why they must be applied to only those actions that under the agent's control. I will argue, in the next section, that there are in fact three roles that deontic principles play in

since evaluative principles are by themselves sufficient to play the role of determining which options are better than others and by how much. (Of course, Norcross is talking about morality and whether there is more *moral* reason to perform one act than to perform another, whereas I'm talking about practical normativity in general and whether there is more reason, *all things considered*, to perform one act than to perform another. Nevertheless, Norcross seems to think that what holds for morality also holds for other practical normative realms, such as prudence. See Norcross 2006, p. 47.) As we'll see below, I think that it is a mistake to regard deontic principles as superfluous. Deontic principles tell us which acts are obligatory. And knowing which acts are obligatory helps us to determine both when certain types of agents are criticizable/blameworthy for their actions and when we can necessarily expect certain types of agents to perform these obligatory actions. For further criticisms of Norcross's view, see Portmore 2011a, pp. 143–145.

addition to their classificatory role and that corresponding to each of these roles there is a condition for the plausibility of any proposed deontic principle.⁹ These conditions will then be used in the following section to assess the relative plausibility of various possible versions of PYBxO.

2. Plausibility Conditions

As noted above, we can use evaluative principles to evaluate almost anything, including the climate, seismic events, a person's eye color, and the annual rainfall in our city. But deontic principles are not like this. Deontic principles can be used to evaluate only that which is, in some relevant sense, under our control. Yet, if the only role that deontic principles had to play was an evaluative one, this restriction would seem to be misplaced. But deontic principles differ from evaluative principles precisely in that they have additional roles to play. One of these additional roles is a practical one. Deontic principles, unlike evaluative principles, must be able to serve as a guide to our rational deliberations about what we should intend to do. And this is why deontic principles are restricted to those things that are under our deliberative control—that is, to those actions that are such that whether or not they are performed depends on our rational deliberations and the intentions that result from them. Whether my toenails or fingernails grow at a faster rate depends neither on my rational deliberations nor on the intentions that I form as a result of them, and this is why it makes no sense to claim that I'm obligated to grow one faster than the other.¹⁰ By contrast, whether I raise my arm depends both on my rational deliberations (specifically, on my deliberations concerning whether to raise my arm) and on the intention that I ultimately form as a result of these deliberations. Deontic principles, unlike

⁹ In enumerating these three roles, I have benefitted from reading Hedden 2011. But whereas Hedden discusses three roles for subjective *oughts*, I discuss the three corresponding roles for objective *oughts*.

¹⁰ Global consequentialists would disagree, but see Chappell *forthcoming* for why they're mistaken. As Chappell notes, although we can rightly criticize someone for, say, *failing to desire* to have a certain eye color (e.g., one that it would be good for her to have), we cannot rightly criticize someone for *failing to have* (or *to come to have*) a certain eye color, simply because what eye color she has (or comes to have) isn't under her rational control in the way that what desires she has is—that is, in being responsive to the her reasons. Likewise, we might rightly criticize someone for failing to desire to grow both their toenails and their fingernails at the exact same rate (supposing that it would be better if this were the case), but we cannot rightly criticize someone for failing to grow both their toenails and their fingernails at the exact same rate.

evaluative principles, serve this practical role of guiding our deliberations about what to intend to do. Evaluative principles can tell us whether we have more reason to desire that we perform this alternative as opposed to that alternative, but they can't tell us whether we ought to intend to perform either of these two alternatives.

Of course, deontic principles come in at least two flavors: subjective and objective. A principle that holds that agents ought to maximize expected utility is subjective, because it makes what an agent ought to do a function of her evidence regarding the relevant facts. By contrast, a principle that holds that agents ought to maximize actual utility is objective, because it makes what an agent ought to do a function of the relevant facts, regardless of what her evidence about them is. Both types of deontic principles are meant to be action-guiding, only for different kinds of agents. Whereas subjective deontic principles are meant to be action-guiding for actual agents with perhaps limited information, objective principles are meant to be action guiding only for those agents (hypothetical or actual) who know all the relevant facts. For my purposes, it will be helpful to have an easy way of distinguishing these two types of agents. So let 'S' range over actual agents, and let 'S+' range over their epistemically ideal counterparts. (Of course, sometimes S and S+ will be the same, as where the actual agent knows all the relevant facts.) Whereas S may or may not know that she ought to comply with a given version of $PYBxO$ and may or may not know which of her x options comply with $PYBxO$, S+ knows these things.

Since, in this paper, I happen to be interested in objective deontic principles, the issue for me will be whether a given version of $PYBxO$ can be used by S+ as a guide to action.¹¹ And for S+ to be able to use a given version of $PYBxO$ as a guide to action is for her to be able to use her knowledge of that principle and what she needs to do to comply with it to ensure that she does comply with it. Suppose, for instance, that according to a given version of $PYBxO$, S has an obligation at t to do ϕ at t' . The idea, then, is that an epistemically ideal agent (S+)—an agent who knows at t both that she ought to comply with this version of $PYBxO$ and that she must ϕ at t' in order to do so—can use this information at t in conjunction with a desire (or other motive) at t to comply with this version of $PYBxO$ to bring about her compliance with this $PYBxO$ -obligation. More precisely, S+ uses $PYBxO$ at t as a guide to fulfilling her $PYBxO$ -obligation at t to ϕ at t' if and only if S+ has a $PYBxO$ -obligation at t to ϕ at t' and comes to ϕ at t'

¹¹ Much of what I say regarding objective principles could be said about subjective principles, *mutatis mutandis*.

out of both a desire (or other motive) at t to comply with $PYBxO$ and a belief at t that she must ϕ at t' in order to comply with $PYBxO$.¹²

Given this practical role that deontic principles are meant to play, we can now state our first plausibility condition for any proposed version of $PYBxO$:

The Practical Condition: A given version of $PYBxO$ is plausible only if it is such that $S+$ can use it at t as a guide to fulfilling whatever $PYBxO$ -obligations that she has at t .

Besides this practical role, deontic principles play a role in the assessment of agents with respect to their actions. Indeed, it is often assumed that there is a conceptual connection between acting impermissibly and being appropriately criticizable for so acting.¹³ Of course, no one thinks that the connection is so simple that if an agent acts so as to violate some objective deontic obligation, then she is appropriately criticizable for so acting. For the agent may not be responsible for her actions—perhaps her actions were coerced, perhaps she lacked the relevant control over her actions, or perhaps she failed to meet some epistemic condition for responsibility. But if we restrict ourselves to epistemically ideal agents (i.e., $S+$) and stipulate that $S+$ responsibly ϕ s at t' if and only if $S+$ ϕ s at t' and is responsible for ϕ -ing at t' such that $S+$ is potentially the appropriate target of praise or criticism for having ϕ -ed at t' (whether it be praise or criticism or neither that's appropriate will depend on ϕ 's deontic status), then we can express the conceptual connection as follows: it is appropriate to criticize $S+$ for responsibly ϕ -ing at t' if and only if $S+$ violates some objective obligation in ϕ -ing at t' .¹⁴ Given this conceptual connection, we should accept the following as our second plausibility condition:

The Assessment Condition: A given version of $PYBxO$ is plausible only if it is such that it is appropriate to criticize $S+$ for responsibly ϕ -ing at t' if and only if $S+$ violates some $PYBxO$ -obligation in ϕ -ing at t' .

¹² I'm borrowing this account of using a deontic principle to guide action, with minor revisions, from H. M. Smith 1988—see especially p. 92.

¹³ See, for instance, Darwall 2006, Gibbard 1990, Mill 1991, Portmore 2011, and Skorupski 1999.

¹⁴ I'm allowing that one can be excused for having failed to fulfill one's obligations if one had some suitable excuse for this failure, but I'm stipulatively defining 'responsibly acting' such that it is conceptually impossible for one to responsibly act in a way that contravenes one's obligations while having some suitable excuse for this failure.

So far, we've seen that objective deontic principles play both a practical role in guiding epistemically and motivationally ideal agents to fulfill their obligations and an assessment role in determining the appropriateness of criticizing epistemically ideal agents for responsibly acting in ways that violate their objective obligations. In addition to these two roles, objective deontic principles also play a predictive (and explanatory) role, allowing us to predict how perfectly rational and epistemically ideal agents will behave and to explain why they behaved as they did. Specifically, we should be able to predict that agents who are both perfectly rational and epistemically ideal (possessing all the relevant information necessary for determining what their objective rational obligations are) will necessarily act in accordance with their objective rational obligations.¹⁵ And this means that we should accept the following as our third plausibility condition:

The Predictive Condition: A given version of PYBxO is plausible only if it is such that we can predict that an S+ who is perfectly rational at t will necessarily ϕ at t' if and only if S+ has a PYBxO-obligation at t to ϕ at t' .

The above three conditions correspond to the three additional roles that are played by objective deontic principles—additional, that is, to the role that both deontic and evaluative principles play in grouping acts into various deontic/evaluative categories. If it weren't for their playing these additional roles, deontic principles would be superfluous, for evaluative principles are perfectly adequate when it comes to classifying actions in terms of how much reason we have to perform them. We can, for instance, use evaluative principles to determine both that ridding the universe of all suffering is better than helping an old lady across the street and that destroying the universe (and all the creatures in it) is worse than squashing a single bug. But knowing which acts are better (or worse) than others is not enough. In addition, it's important to know what we should do, when it's appropriate to criticize someone for responsibly acting in a certain way, and whether we can necessarily predict that someone will behave in a certain way insofar as he or she is perfectly rational and ideally informed. Certainly, it would be good if I were to rid the universe of all suffering, but am I

¹⁵ I take it that, as a matter of conceptual necessity, a perfectly rational and epistemically ideal agent will always do what she has decisive reason to do and will only do what she has sufficient reason to do. And I'm assuming both that S has decisive reason to ϕ if and only if S's reasons are such as to make S objectively and all-things-considered required to ϕ and that S has sufficient reason to ϕ if and only if S's reasons are such as to make S objectively and all-things-considered permitted to ϕ .

obligated to do so? Would it be appropriate to criticize me if I failed to do so? And can we predict that I will do so insofar as I am perfectly rational and ideally informed? It's clear that evaluative principles are not up to the task of answering such questions. The only principles that are up to this task are those that are neither too broad nor too narrow to meet the above three conditions.

But before proceeding to discuss various proposed accounts of the relevant options and the corresponding versions of $PYBxO$, it will be helpful to illustrate both the vice of being too broad and the vice of being too narrow, using some clear-cut offenders. Take, first, the vice of being too broad and one obvious culprit: $PYBLPO$ (*perform your best logically possible option*). To see that $PYBLPO$ is clearly too broad, consider the following case in which I am the agent. And let's assume that I am ideally informed (and, thus, an instance of $S+$), perfectly rational, and sufficiently motivated to comply with $PYBLPO$. According to $PYBLPO$, I am obligated now, at t , to rid the universe of all suffering, for it is logically possible for me to do so and this is the best, or so we'll assume, of all my logically possible options. But note that I could not use $PYBLPO$ at t as a guide to fulfilling this obligation. Even though I am sufficiently motivated at t to comply with $PYBLPO$, I cannot use the knowledge that I have of this $PYBLPO$ -obligation to fulfill it. Possessing this motive and knowledge at t is insufficient to bring me to fulfill this obligation, because no matter how much I want (and regardless of whether or not I intend) at t to fulfill this obligation, I won't fulfill it. Thus, $PYBLPO$ fails the practical condition.

$PYBLPO$ also fails the assessment condition, for this principle is not such that it is appropriate to criticize $S+$ for responsibly ϕ -ing at t' if $S+$ violates some $PYBxO$ -obligation in ϕ -ing at t' . Suppose, for instance, that instead of ridding the universe of all suffering, I responsibly act so as to rid the universe of as much suffering as I can. But because the best I can do falls far short of the best that it is logically possible for me to do, I end up violating my $PYBLPO$ -obligation to rid the universe of all suffering. Nevertheless, it would be inappropriate to criticize me for having responsibly acted so as to rid the universe of as much suffering as I can. Thus, $PYBLPO$ fails the assessment condition.

Lastly, $PYBLPO$ fails the predictive condition, for it is not such that we can predict that an $S+$ who is perfectly rational at t will necessarily ϕ at t' if $S+$ has a $PYBxO$ -obligation at t to ϕ at t' . We cannot, for instance, predict that I will rid the universe of all suffering at t' even though I have a $PYBLPO$ -obligation to do so. Indeed, we can predict that I won't in spite of my being perfectly rational, ideally informed, and ideally motivated.

The other vice is that of being too narrow. To illustrate this vice, consider the strange view that you ought to *perform your best non-exertive option* ($PYBNEO$ —pahy-bi-NEE-oh), where a non-exertive option is just an action that requires no

strenuous effort. PYBNEO's account of options is clearly too narrow. Consequently, it fails both the assessment condition and the predictive condition. It fails the assessment condition, because it is not such that it is appropriate to criticize S+ for responsibly ϕ -ing at t' *only if* S+ violates some PYBxO-obligation in ϕ -ing at t' .¹⁶ To see this, consider a variation on *Lake* and assume that, in this version (*Lake 2*), Abe will be able to reach Fred in time by swimming across the lake. Of course, swimming across the lake would require some strenuous effort. Thus, instead of doing this, Abe performs his best non-exertive option, which is to search the internet for a new financier to replace his soon-to-be-dead financier. As a result, his friend dies and he loses valuable time in having to find a new financier for his business. Surely, it is appropriate to criticize Abe for responsibly searching the internet for a new financier rather than swimming across the lake to save his friend and financier. Yet, he violates no PYBNEO-obligation in choosing to search the internet. So PYBNEO fails the assessment condition.

PYBNEO also fails the predictive condition, for it is not such that we can predict that an S+ who is perfectly rational at t will necessarily ϕ at t' *only if* S+ has a PYBxO-obligation at t to ϕ at t' .¹⁷ Indeed, we can clearly predict that Abe+ will, if perfectly rational, necessarily swim across the lake even though he has no PYBNEO-obligation to do so. After all, it would be positively foolish for him not to swim across the lake and save his friend and financier when he could do so with only a little strenuous effort. Besides, swimming across the lake would give Abe some much needed exercise. So insofar as Abe is perfectly rational and ideally informed, we can predict that he will necessarily violate PYBNEO. Thus, PYBNEO clearly fails the predictive condition as well as the assessment condition.

3. Options

Below, I consider various proposals for what we might substitute for x in PYBxO. I argue that some are too broad and that others are too narrow. In the end, I find a substitution for x that seems just right: ' x ' equals 'scrupulously securable'. But,

¹⁶ When a principle is too broad, the problem, by contrast, is that it is not such that it is appropriate to criticize S+ for responsibly ϕ -ing at t' *if* S+ violates some PYBxO-obligation in ϕ -ing at t' .

¹⁷ When a principle is too broad, the problem, by contrast, is that it is not such that we can predict that an S+ who is perfectly rational at t will necessarily ϕ at t' *if* S+ has a PYBxO-obligation at t to ϕ at t' .

to begin with, consider what is perhaps the first thing that comes to mind when looking for a plausible candidate for x : ' x ' equals 'effectively decidable', where

EDO S 's ϕ -ing at t' is an *effectively decidable option* for S at t if and only if S would ϕ at t' if S were to intend at t to ϕ at t' ($t < t'$).¹⁸

The resulting view, which holds that you ought to *perform your best effectively decidable option* is PYBEDO (pahy-BEE-doh). PYBEDO is too broad. To see this, consider *Businesswoman*. In order to save her business, Jones (a businesswoman) needs to quickly transfer some funds from one account to another. And let's assume that Jones is ideally informed (and, thus, an instance of $S+$), perfectly rational, and sufficiently motivated to comply with PYBEDO. Being ideally informed, Jones knows that it's crucial for her to transfer the funds, and, being perfectly rational, she is just about to complete the transfer online by clicking on the transfer button, when she is shot in the head by her business rival, rendering her unconscious. So Jones is now, at t_1 , lying unconscious on the floor beside her computer.¹⁹ Nevertheless, according to PYBEDO, she is obligated now, at t_1 , to click on the transfer button at t_2 , for this is her best effectively decidable option at t_1 . Clicking on the transfer button at t_2 is effectively decidable at t_1 , because, if she were to intend at t_1 to click on the transfer button at t_2 , she would succeed in

¹⁸ Goldman proposes this view—see her Definition II (1978, p. 195). In note 34 (p. 213), though, she concedes that this definition is probably only a rough approximation of the correct definition.

Also, note that ' ϕ ' will sometimes stand for 'try to x ' as opposed to ' x ', where ' x ' stands for something difficult to do such as 'hit a hole-in-one'. After all, I would never intend to hit a hole-in-one, for that would be presumptuous. Instead, I would intend only to *try* to hit a hole-in-one (Mele 1989). And although it may not be determinate whether I would hit a hole-in-one if I were to try to do so, it can be determinate whether I would try to hit a hole-in-one if I were to intend to try to hit a hole-in-one. It seems that intending to ϕ entails believing that one will ϕ (Harman 1976 and Velleman 1989) and that this explains both why I cannot rationally intend to hit a hole-in-one when I don't have much credence in the thought that I will and why we should think that subjunctive conditionals such as " S would try to hit a hole-in-one at t' if S were to intend at t to try to hit a hole-in-one at t'' " will have determinate truth values even if subjunctive conditionals such as " S would hit a hole-in-one at t' if S were to intend at t to try to hit a hole-in-one at t'' " don't. (See Hare 2011 and Vessel 2003 for more on the problem of underspecified conditional antecedents.)

¹⁹ Assume that although she is unconscious, she is in some dispositional sense ideally informed, perfectly rational, and properly motivated—that is, in the sense that someone could be said to possess the relevant dispositional beliefs, dispositional desires, and other non-occurrent mental states even while under general anesthetic.

doing so. This is true in virtue of the fact that the closest possible world in which she intends at t_1 to click on the transfer button at t_2 (let's say the one in which the bullet takes a slightly different path through her brain, consequently leaving her conscious throughout t_1 and t_2) is the one in which she is conscious throughout t_1 and t_2 and is, thus, forms and carries out an intention to click on the transfer button at t_2 .²⁰ But the idea that she has at t_1 , while lying unconscious on the floor with no chance of regaining consciousness anytime soon, the property of being obligated to click on the transfer button at t_2 is quite implausible.

Part of the reason that PYBEDO is so implausible is that it fails the practical condition.²¹ There is no way for Jones to use PYBEDO at t_1 as a guide to fulfilling her PYBEDO-obligation at t_1 to click on the transfer button at t_2 . To be able to use PYBEDO as a guide to fulfilling this obligation, Jones would have to be such that she could come to click on the transfer button at t_2 out of both a desire (or other motive) at t_1 to comply with PYBEDO and a belief at t_1 that she must click on the transfer button at t_2 to comply with PYBEDO.²² Thus, she must be able at t_1 to form an intention to click on the transfer button at t_2 as a result of possessing both this belief and this desire. But, given her unconscious state, Jones lacks the capacity to form any intention. Thus, PYBEDO fails the practical condition.

In light of this failure, it may seem that all we need to do is to modify EDO so that S 's having at t the capacity to continue, or to come, to have at t' the intention to ϕ at t'' is a necessary condition for S 's having ϕ -ing at t'' as an x option at t . In other words, we might think that all we need to do is to substitute 'doable' instead 'effectively decidable' for ' x ', where

DO S 's ϕ -ing at t'' is a *doable option* for S at t if and only if both (a) S would ϕ at t'' if S were to intend at t' to ϕ at t'' and (b) S has at t the capacity to continue, or to come, to have at t' the intention to ϕ at t'' ($t \leq t' < t''$).²³

²⁰ I'm assuming that we should adopt something along the lines of the Lewis-Stalnaker approach to counterfactuals.

²¹ Clearly, PYBEDO also fails both the assessment condition and the predictive condition.

²² I'm assuming that, in some dispositional sense, Jones has at t both this desire and this belief.

²³ It's difficult to spell out exactly what constitutes having such a capacity, but I think that this much is obvious: S has at t the capacity to continue, or to come, to have at t' the intention to ϕ at t'' only if S has at t rational control over whether or not she continues, or comes, to have at t' the intention to ϕ at t'' . And, for her to have this sort of control, I presume that she must be at t conscious, capable of deliberating, and in possession of the necessary conceptual apparatus. I leave open, though, what else might be required, including whether S 's not being causally determined to continue, or to come, to have at t'

The resulting theory is PYBDO (PAHY-bi-doh), but unfortunately it is too narrow. To see why, consider:

Foolishness: Phyllis is a *physician* and a recovering alcoholic. Her *patient*, Patience, is suffering from a very painful and potentially life-threatening medical condition. If Patience receives one dose of Drug C ('C' for *completely* cured) at 1 P.M. Monday and a second dose of Drug C at 1 P.M. Tuesday, she will be completely cured and both she and Phyllis will live happily ever after. If Patience receives one dose of Drug P ('P' for *partially* cured) at 1 P.M. Monday and a second dose of Drug P at 1 P.M. Tuesday, she will be partially cured, ensuring that she won't die but also ensuring that she'll spend the rest of her life in pain.²⁴ If this happens, Phyllis will lose her job and be sued for malpractice. If Patience is given any other treatment besides these two or is left untreated, she'll die and Phyllis will spend the rest of her life in prison for negligent homicide. Now, Phyllis is the only person who can administer Drug C. And unless Phyllis arrives at the hospital by 1 P.M. Tuesday, the hospital staff will, given standing orders, administer Drug P to Patience at 1 P.M. Tuesday. It's just before 1 P.M. Monday and Phyllis must decide what to do. One thing that she must keep in mind is that Tuesday is the anniversary of her son's death, which is what drove her to alcoholism in the first place. Thus, she will be sorely tempted to drink on Tuesday morning and, if she does, she'll end up passed out drunk by noon on Tuesday and the hospital staff will have no choice but to administer Drug P to Patience at 1 P.M. Tuesday.

Ideally, Phyllis administers Drug C at 1 P.M. Monday, meets up with her Alcoholics Anonymous sponsor first thing Tuesday morning, refrains from drinking all Tuesday with the help of her sponsor, and administers the second dose of Drug C at 1 P.M. Tuesday. And so long as she intends now (i.e., just before 1 P.M. Monday) do all of these things, she will succeed in doing them. Unfortunately, though, Phyllis has no intention at present of

some other incompatible intention is also necessary for *S*'s having at *t* the capacity to continue, or to come, to have at *t'* the intention to ϕ at *t''*. Note that *S* has at *t* rational control over whether or not she continues, or comes, to ψ only if she is at *t* at least moderately responsive to the relevant types of reasons. See Fischer & Ravizza (1998, pp. 62–91 and especially 243–244) for an account of moderate reasons-responsiveness.

For a helpful discussion of views such as EDO and DO, from which my discussion draws heavily, see Carlson 1995, pp. 77–82.

²⁴ Assume that the pain is such that she would be better off living with that pain than dying, but not by much.

meeting up with her sponsor Tuesday morning. Although she recognizes that she has a very good reason to do so (specifically, that this would ensure that she won't give into the temptation to drink on Tuesday), she fails to respond appropriately to this reason. And so instead of planning to meet up with her sponsor on Tuesday (as she should), she plans to spend the day alone. Given her plan/intention to spend Tuesday alone, she's going follow through with this intention. And her spending the day alone will result in her giving into the temptation to drink. Consequently, she is going to end up passed out drunk and unable to administer a second dose of Drug C on Tuesday. (Assume that Phyllis is an instance of S+ and, thus, knows all the relevant facts.)²⁵

According to PYBDO, Phyllis is under no present obligation to completely cure her patient (that is, to administer one dose of Drug C at 1 P.M. Monday and another dose of Drug C at 1 P.M. Tuesday), for she can be obligated at present to do only that which is a doable option for her at present. And her administering a second dose of Drug C on Tuesday is not a doable option for her at present. It's not a doable option, because condition *a* of DO is not met. Given that Phyllis at present intends to spend Tuesday alone, it is not the case that Phyllis would administer a second dose of Drug C Tuesday even if she were to intend now to do so. For given her lack of any intention to meet up with her sponsor on Tuesday, her present intention to administer the second dose of Drug C on Tuesday will prove ineffective. Instead of carrying through with her intention to administer a second dose of Drug C on Tuesday, she is going to end up passed out drunk.

Now, suppose that Monday and Tuesday have already passed and that the following has taken place. Just before 1 P.M. Monday, Phyllis decided to give Patience a dose of Drug P at 1 P.M. Monday. She decided not to administer Drug C on Monday, for she reasoned that even if she were to intend now to administer Drug C on Tuesday, she would not follow through with this intention. As a result of this reasoning and her resulting decision, she administered Drug P at 1 P.M. Monday, spent Tuesday morning alone getting drunk, and then ended up passed out drunk by noon, leaving the hospital staff to administer a second dose of Drug P at 1 P.M. Tuesday. Patience has, then, received only a partial cure, and although she won't die, she will suffer from chronic, incurable pain for the rest of her life. And Phyllis has lost her job and is being sued for malpractice.

Can we appropriately criticize Phyllis for having failed to administer Drug C on both Monday and Tuesday? It seems that we can. After all, Phyllis recognized

²⁵ This example is based on a similar one found in Feldman 1986, p. 11.

on Monday that she had a good reason to meet up with her sponsor on Tuesday morning. She just failed to respond appropriately to this reason and, thus, failed to form the intention to meet up with her sponsor. And this is so despite her having, we'll assume, the capacity to respond appropriately to such reasons. So, when she was deliberating about what to do just before 1 P.M. Monday, she should not have held fixed, as she did, the ineffectiveness of any present intention to administer Drug C on Tuesday. For the only reason that this intention was to be ineffective was because she wasn't in her present rational deliberations responding appropriately to her reasons for intending to meet up with her sponsor on Tuesday morning. And agents should not, while rationally deliberating, treat as fixed that which depends on their rational deliberations.

Oftentimes, whether or not an agent's present action will achieve its intended purpose depends on whether or not it will be followed up by some future action, an action performed either by the agent's future self or by some other future agent. Such is the case with Phyllis. Administering a dose of Drug C at 1 P.M. today (Monday) will help to cure Patience only if Phyllis follows up with a second dose of Drug C at 1 P.M. tomorrow (Tuesday). If she doesn't follow up with a second dose of Drug C tomorrow, then not only will Patience not be cured, but she'll die. In these types of cases, there are two mutually exclusive stances that the agent can take towards the relevant future action: the predictive stance and the deliberative stance. When an agent adopts the predictive stance towards some future action, she must try to predict whether it will occur. When, by contrast, she adopts the deliberative stance towards some behavior, she must put aside any predictions she has about what she will do and instead decide what she will do.²⁶

Phyllis, unfortunately, adopted on Monday the predictive stance towards her administering a second dose of Drug C on Tuesday, and she predicted that she would not administer a second dose of Drug C. She rightly predicted that, given her lack of any intention to meet up with her sponsor on Tuesday, she was going to end up passed out drunk on Tuesday, unable to administer a second dose of Drug C. Yet it seems that she should have adopted the deliberative stance towards her administering a second dose of Drug C on Tuesday. For, at the time, whether or not she was going to do so depended on her present deliberations. If, instead of concluding in her deliberations that she should administer a dose of

²⁶ I borrow the distinction between the predictive stance and the deliberative stance from Woodard 2008. However, I explicate the distinction slightly differently than he does. And also I disagree with him on whether it is ever appropriate to adopt the deliberative stance towards behavior that is not under one's present deliberative control. He thinks that it is, and I disagree.

Drug P on Monday because she was going to end up passed out drunk on Tuesday, she had concluded that she should meet up with her sponsor on Tuesday and administer doses of Drug C on both Monday and Tuesday, then that's what she would have done. Thus, her administering a dose of Drug C on Tuesday was, on Monday, under her deliberative control—that is, whether or not she was to perform this future act depended entirely on the course and conclusions of her deliberations on Monday. Given this, it seems inappropriate for her to have adopted the predictive stance as opposed to the deliberative stance towards her behavior on Tuesday. Agents should adopt the predictive stance towards only that behavior that is not under their present deliberative control and should adopt the deliberative stance towards all behavior that is under their present deliberative control.

Given that Phyllis failed on Monday to respond appropriately to her reasons and failed on Monday to adopt the appropriate stance to her actions on Tuesday, we can rightly criticize her for not having administered Drug C on both Monday and Tuesday. For had she adopted the deliberative stance towards her future actions and responded appropriately to her reasons while she was deliberating, she would have formed both the intention to meet up with her sponsor on Tuesday morning and the intention to administer Drug C on both Monday and Tuesday and would have, thereby, succeeded in doing all those things.

So it seems that PYBDO fails the assessment condition, for PYBDO is not such that it is appropriate to criticize S+ for responsibly ϕ -ing at t' only if S+ violates some PYBDO-obligation in ϕ -ing at t' . After all, it is appropriate to criticize Phyllis for (responsibly) failing to completely cure Patience even though Phyllis did not thereby violate any PYBDO-obligation. The lesson seems to be that even if an agent cannot ensure that she'll perform some future action simply by intending now to perform that future action, she may be able to ensure that she'll perform that future action by intending now to perform some set of actions that includes both that future action and whatever other acts she ought now to intend to perform. And, in that case, she is not excused from criticism for failing to perform that future action simply because she didn't intend to perform the other actions that she should have intended to perform and that would have rendered her intention to perform that future action effective. This suggests that we should substitute, not 'doable', but 'securable' for ' x ', where

- SO** S's ϕ -ing at t'' is a *securable option* for S at t if and only if there is a time, t' (where t' is either identical or immediately subsequent to t), and a set of actions, ψ , (where ψ may, or may not, be identical to ϕ) such that both of the following are true: (1) S would ϕ at t'' if S were to

have at t' the intention to ψ and (2) S has at t the capacity to continue, or to come, to have at t' the intention to ψ ($t \leq t' < t''$).²⁷

Although the resulting view, PYBSO (PAHY-bi-soh), gets the right answer in *Foolishness*, implying that Phyllis is obligated just before 1 P.M. Monday to administer a dose of Drug C on both Monday and Tuesday, it is still too narrow. For if we think that an agent can't avoid incurring an obligation at t to ϕ at t'' simply because she fails at t to respond appropriately to her reasons, thereby failing to come to have certain *intentions* that she ought to have and that she would need to have for her intention to ϕ at t'' to be effective, then we should also think that she can't avoid incurring an obligation at t to ϕ at t'' simply because she fails at t to respond appropriately to her reasons, thereby failing to come to have certain *beliefs* (or other judgment-sensitive attitudes) that she ought to have and that she would need to have for her intention to ϕ at t'' to be effective.²⁸ After all, there is no relevant difference between intentions and other judgment-sensitive attitudes, such as beliefs and desires. Neither intentions nor other judgment-sensitive attitudes are under our volitional control, and both are under our rational control.²⁹

To illustrate the problem for PYBSO, consider the following case:

²⁷ I borrow the term 'securable' and the basic notion of securability from Sobel 1976, p. 96. I do, however, formulate things a bit differently. For me, what's relevant is whether there is some intention that will secure S's ϕ -ing at t'' . For Sobel, what's relevant is whether there is some minimal act (i.e., an act that once begun cannot be stopped short of its completion) open to S whose present performance would secure S's ϕ -ing at t'' . But, perhaps, there is no minimal act that will itself secure S's ϕ -ing at t'' , but there is some intention that will, regardless of what minimal act S presently performs, secure S's ϕ -ing at t'' . In that case, S's ϕ -ing at t'' will be securable on my view, but not on Sobel's. Suppose, for instance, that I'll go for a walk this evening if and only if I now intend to do so, and assume that I do not now intend to do so. In that case, there would be no minimal act that, if presently performed by me, would secure my going for walk this evening, for it's an intention, not some act, that secures my going on a walk this evening. Thus, in this case, my going for a walk this evening is securable on my view, but not on Sobel's view.

²⁸ Judgment-sensitive attitudes—such as, beliefs, desires, and intentions—are attitudes that are sensitive to our judgments about reasons for and against having them. They contrast with other mental states—such as, hunger—that are insensitive to such judgments. See Scanlon 1998, p. 20.

²⁹ Recall that S has rational control over whether or not she ϕ s only if both (1) she has the capacity to recognize and assess the relevant types of reasons and (2) her ϕ -ing is at least moderately responsive to her judgments concerning these types of reasons. By contrast, S has volitional control over whether or not she ϕ s only if both (1) she has the capacity to intend to ϕ and (2) whether or not she ϕ s depends on whether or not she intends to ϕ .

Wishful Thinking: This case is like *Foolishness*, except that in this case the problem is not that Phyllis doesn't intend to meet up with her sponsor on Tuesday morning. In fact, she does. The problem is that Phyllis has failed to respond appropriately to her evidence and has, as a result of wishful thinking, formed the false belief that she can resist the temptation to drink without meeting up with her sponsor. And, given this belief, she will change her mind about meeting up with her sponsor on Tuesday morning. For, on Tuesday morning, she'll find that she prefers to spend the day alone, and, thinking that she can resist the temptation to drink on her own, she'll change her mind and decide to spend the day alone. And, as a result, she'll end up passed out drunk and unable to administer the second dose of Drug C on Tuesday. (Assume that Phyllis is an instance of S+.³⁰)

According to PYBSO, Phyllis is under no present obligation to administer a second dose of Drug C on Tuesday, for she can be obligated at present to do only that which is a securable option for her at present. And her administering a dose of Drug C on Tuesday is not a securable option for her at present. This is because condition 1 of SO is not satisfied. Given that Phyllis presently believes that she can resist the temptation to drink on Tuesday all on her own, it is not the case that she would administer a dose of Drug C on Tuesday even if she were now to intend both to meet up with her sponsor Tuesday morning and to administer a dose of Drug C on Tuesday. Indeed, no matter what set of actions she intends now to perform, she will not administer a dose of Drug C on Tuesday. For, given her belief that she can resist the temptation to drink all on her own, her present intention to meet up with her sponsor is quite precarious and will in fact be abandoned in the face of a later desire to spend the day alone. And her spending the day alone will be disastrous both for her and for Patience.

Now, suppose that Monday and Tuesday have already passed and that the following has taken place. Just before 1 P.M. Monday, Phyllis decided she would give Patience the first dose of Drug C at 1 P.M., meet up with her sponsor tomorrow (Tuesday) morning, and administer the second dose of Drug C at 1 P.M. Tuesday. At 1 P.M. Monday, Phyllis did in fact administer the first dose of Drug C. But, on Tuesday morning, she changed her mind about meeting up with

³⁰ The fact that Phyllis is an instance S+ doesn't entail that she knows (and, thus, believes) that she can resist the temptation to drink without meeting up with her sponsor. That would be contrary to the stipulations of the case. Rather an agent is an instance of S+ if and only if she knows that she ought to comply with the given version of PYBxO (in this case, PYBSO) and also knows which of her *x* options comply with PYBxO.

her sponsor and decided to spend the day alone. Later that morning while agonizing over her son's death, she was overcome by the temptation to drink. She then drank until she passed out. Consequently, she never arrived at the hospital that afternoon. Thus, the hospital staff had no choice but to administer Drug P. As a result, Patience died, and Phyllis is now serving a life sentence for negligent homicide.

Can we appropriately criticize Phyllis for having failed to administer the second dose of Drug C on Tuesday? It seems that we can. After all, Phyllis had overwhelming evidence to suggest that she wouldn't be able to resist the temptation on her own: indeed, the last two times she tried, she failed. But she just ignored this evidence in her deliberations and, through wishful thinking, formed the belief that she could resist the temptation on her own. And she did this despite having, as we'll presume, the capacity to respond appropriately to her evidence. In light of this, I think that we can rightly criticize Phyllis for not having administered Drug C on Tuesday. For we know that, had she responded appropriately to her reasons when she was deliberating just before 1 P.M. Monday, she would have formed all the following attitudes: (1) the belief that meeting up with her sponsor is essential to her resisting the temptation to drink, (2) the intention to meet up with her sponsor on Tuesday morning, and (3) the intention to administer Drug C on both Monday and Tuesday. And had she formed this combination of attitudes (as she ought to have), she would have succeeded in following through with her intention to administer Drug C on both Monday and Tuesday.

So it seems that PYBSO fails the assessment condition, for PYBSO is not such that it is appropriate to criticize S+ for responsibly ϕ -ing at t' only if S+ violates some PYBSO-obligation in ϕ -ing at t' . After all, it is appropriate to criticize Phyllis for having responsibly failed to administer a second dose of Drug C at 1 P.M. Tuesday even though Phyllis did not thereby violate any PYBSO-obligation. The lesson seems to be that even if an agent cannot, while maintaining her current set of background attitudes, ensure that she'll perform some future action by intending now to perform it or some other suitable set of actions, she may be able to ensure that she'll perform that future action so long as she responds appropriately to her reasons and, thus, comes to have all the judgment-sensitive attitudes that she ought to have. If so, she is not excused from criticism for failing to perform that future action simply because she didn't have certain judgment-sensitive attitudes that she should have had and that, had she possessed, would have ensured that she performed that future action. This suggests that what's relevant is what one can secure through one's present intentions while having the set of background attitudes that one ought to have. Thus, instead of 'securable', we should substitute 'scrupulously securable' for ' x ', where:

SSO S 's ϕ -ing at t'' is a *scrupulously securable option* for S at t if and only if there is a time t' (where t' is either identical to or immediately subsequent to t), a set of intentional objects ψ (where ψ may, or may not, be identical to ϕ), and a set of background attitudes B such that all of the following are true: (1) S would ϕ at t'' if S were to have at t' both B and the intention to ψ ; (2) S has at t the capacity to continue, or to come, to have at t' both B and the intention to ψ ; and (3) S would at t' continue, or come, to have B (and, where ψ is not identical to ϕ , the intention to ψ as well) if S were at t aware of all the relevant reason-constituting facts and were to respond appropriately to them, thereby coming to have at t' all those attitudes that, given those facts, she has decisive reason to have and only those attitudes that she has, given those facts, sufficient reason to have ($t \leq t' < t''$).³¹

Whereas the preceding views seemed either too broad or too narrow, PYBSSO seems just right. Or so it seems to me. But others will disagree and argue that it is too narrow and, thus, gets the wrong answer in the following case.

Temptation: This case is somewhat like *Foolishness*, except crucially that in this case Phyllis cannot at present scrupulously secure her administration of a

³¹ After I presented this view in Portmore 2011a, I subsequently discovered Sobel 1982. Both Sobel (1976) and I originally advocated something along the lines of PYBSO. And both of us subsequently revised our original views in light of subsequent objections so as to come to endorse PYBSSO or something very close to it. Sobel (1982) was responding to objections raised by Bergström (1977), whereas I (2011a) was responding to objections that Jacob Ross presented in an email discussion that we had in 2010. In Sobel 1982, Sobel presents something very similar to PYBSSO—see pp. 118–119. His view differs slightly from mine, though. On his view, S 's ϕ -ing at t'' is a scrupulously securable option (or what he calls an “ideally securable” option) for S at t if and only if S 's ϕ -ing at t'' would be securable at t if S “were at t free of all mistakes of heart and mind.” So, on his view, it is sufficient that S 's ϕ -ing at t'' would be securable for S at t if S were at t to respond appropriately to her reasons for having various intentions and other judgment-sensitive attitudes. But, on my view, this is not sufficient, for, given condition 2 of SSO, S must have the capacity to respond appropriately to her reasons in that she must at least be moderately reasons-responsive—see note 23 above. So there could be cases—cases where the agent is guilty of certain present mistakes that she lacks the capacity to rectify—where an act that's scrupulously securable on Sobel's view is not scrupulously securable on my view.

dose of Drug C on Tuesday. The problem is that no matter what her intentions are now and, indeed, no matter how resolute she is now in her intention to meet up with her sponsor on Tuesday morning, she's going to give into the temptation to start drinking early Tuesday morning. More precisely, no matter what her intentions are now, she won't show up at the hospital on Tuesday to administer a dose of Drug C so long as she has the background attitudes that she ought to have. Of course, if she were, come Tuesday morning, to intend to meet up with her sponsor, she would then meet up with her sponsor, refrain from drinking, and show up at the hospital in time to administer the second dose of Drug C. And having arrived at the hospital, she would then administer a second dose of Drug C at 1 p.m. Tuesday if she were to intend to do so when she arrives. So there is a schedule of intentions over the next twenty-four plus hours such that, if Phyllis's intentions were to follow this schedule, she would carry out all the intentions in that schedule and consequently administer two doses of Drug C at the appropriate times. It's just that, as a matter of fact, her intentions are not going to follow this schedule, for she's not going to intend to meet up with her sponsor come Tuesday morning. She will instead give into temptation and form the intention to start drinking right away. And this is so no matter what intentions she has now.

On PYBSSO, Phyllis has no obligation at present to administer a second dose of Drug C at 1 P.M. Tuesday, for this is not a scrupulously securable option for Phyllis at present. Absent her forming some irrational attitude (such as the belief that her son will come back to life if she refrains from drinking on Tuesday), her present intention to meet up with her sponsor on Tuesday will not be carried out, as she will in fact change her mind about meeting up with her sponsor come Tuesday morning. Indeed, no matter what she does or intends to do now, she will not administer a dose of Drug C tomorrow. The only thing that might change this is if Phyllis's somehow forms an irrational attitude that would then render her present intention less susceptible to abandonment in the face of future temptation—e.g., the belief that her refraining from drinking on Tuesday will raise her son from the dead. But since such an attitude is clearly irrational and, thus, a failure to meet condition 3 of SSO, Phyllis's administering a second dose of Drug C on Tuesday is not a scrupulously securable option. Thus, according to PYBSSO, Phyllis avoids a present obligation to completely cure Patience in virtue of her looming future failure (viz., her future failure to resist the temptation to drink on Tuesday morning). And some find it counterintuitive to think that an agent can avoid a present obligation in virtue of her own looming future failure.

I don't find this counterintuitive at all. Of course, I admit that an agent cannot always avoid a present obligation to perform a future action simply in virtue of some looming future failure. When, for instance, that future failure could be avoided by responding appropriately at present to one's reasons, thereby coming to have all those attitudes that one has decisive reason to have and only those attitudes that one has sufficient reason to have, one cannot avoid such an obligation. However, when that future failure is unavoidable so long as one does at present respond appropriately to one's reasons, then it seems to me entirely plausible to take that looming (and presently unavoidable) future failure as grounds for limiting one's *present* obligations. Such is the case in *Temptation*. If Phyllis does at present respond appropriately to her reasons, she will not administer a dose of Drug C on Tuesday—that's just a stipulation of the case. So, with respect to her present deliberations, she ought to hold fix the fact that she won't administer Drug C on Tuesday and, thus, should administer Drug P on Monday. For her best scrupulously securable option at present is the one in which she administers Drug P on Monday and the staff administers a second dose of Drug P on Tuesday while she's passed out drunk. So I don't see a problem for PYBSSO here. Given that, in *Temptation*, whether or not she will drink on Tuesday morning is not under her present deliberative control, it seems entirely appropriate for her to adopt on Monday the predictive stance towards her future behavior and, thus, to decide to administer Drug P on Monday given her prediction that she will drink on Tuesday and become unable to administer Drug C on Tuesday.

But, in any case, it's worth considering what others might propose as a replacement for PYBSSO. And what they suggest is that we substitute 'personally possible' (or what I'll more concisely refer to as 'performable') for 'x', where:

- PO** *S's* ϕ -ing at t' is a *performable option* for *S* at t if and only if there is some schedule of intentions, I , extending over time-interval, T , beginning at t such that the following are all true: (i) if *S's* intentions followed schedule I , then *S* would carry out all the intentions in I , (ii) *S's* carrying out all the intentions in I would involve *S's* ϕ -ing at t' , (iii) *S* has the capacity at t to have the intentions that I specifies for t , and (iv) for any time t'' in T after t , if *S's* intentions followed I up until t'' , then *S* would have at t'' the capacity to have the intentions I specifies for t'' ($t < t' \leq t''$).³²

³² This formulation draws heavily from the one given in Ross *forthcoming*.

But it seems to me that the resulting view, PYBPO (PAHY-bip-oh), is too broad. To illustrate, consider again *Temptation*. And let's assume that Phyllis is ideally informed (and, thus, is an instance of S+), perfectly rational, and sufficiently motivated to comply with PYBPO. According to PYBPO, Phyllis is obligated at present (just before 1 p.m. on Monday) to administer two doses of Drug C (one today, and one tomorrow). But note that Phyllis could not use PYBPO at present as a guide to fulfilling this obligation. Even though Phyllis wants at present to comply with PYBPO, she cannot use the knowledge that she has a PYBPO-obligation at present to administer two doses of Drug C to fulfill this obligation. Possessing at present both this motive and this knowledge is insufficient to bring about her fulfillment of the obligation, because no matter how much she wants or intends at present to fulfill this obligation, she won't fulfill it. Thus, PYBPO fails the practical condition.

PYBPO also fails the assessment condition, for this principle is not such that it is appropriate to criticize S+ for responsibly ϕ -ing at t' if S+ violates some PYBxO-obligation in ϕ -ing at t' . Suppose, for instance, that instead of two doses of Drug C, Phyllis ensures that Patience gets two doses of Drug P in *Temptation*. I don't see how we can criticize her for failing to fulfill some obligation that she putatively had on Monday to administer two doses of Drug C. After all, on Monday, she knew that the only way for her to save her patient's life and prevent her future incarceration was to administer Drug P on Monday, and that's what she did. She knew that, no matter what she did or intended to do on Monday, Patience was going to get a dose of Drug P (not Drug C) on Tuesday.³³ So how can we blame her for doing the only thing that would prevent both Patience's death and her own life-long incarceration? I don't see how, then, she failed to fulfill any obligation that she putatively had at 1 p.m. on Monday.

Of course, if she were to administer Drug C on Monday, then come Tuesday morning she would have an obligation to meet up with her sponsor and administer the second dose of Drug C, because come Tuesday morning her doing these things would then be scrupulously securable options for her. But this doesn't mean that in failing to meet up with her sponsor on Tuesday morning and administer a second dose of Drug C later that day, she violates any obligation that she had on Monday. Given that she had on Monday no more deliberative control over whether her future self was going to meet up with her sponsor than she had over whether some stranger in some remote and

³³ This is true unless, of course, she were to form some irrational belief, such as the belief that her refraining from drinking on Tuesday would raise her son from the dead. But she's not going to form any such belief, and she couldn't form such a belief if she were to respond appropriately to her reasons.

unreachable location halfway around the world was going to meet up with his sponsor, it seems odd to insist that she had an obligation on Monday to see to it that she meets up with her sponsor on Tuesday but no obligation on Monday to see to it that the stranger meets up with his sponsor on Tuesday. Presumably, given the lack of any deliberative control in each case, she would be right to adopt the predictive stance towards the behavior of each.

Lastly, PYBPO fails the predictive condition, for it is not such that we can predict that an S+ who is perfectly rational at t will necessarily ϕ at t' if S+ has a $\text{PYB}\times\text{O}$ -obligation at t to ϕ at t' . We cannot, for instance, predict that Phyllis will give two doses of Drug C given that she is, as I've stipulated, perfectly rational at present. Indeed, we can predict that, in virtue of her being perfectly rational, she necessarily won't. For what rational person would administer Drug C on Monday, knowing that this would inevitably result in Patience's death and her own life-long incarceration?

Of course, advocates of PYBPO can stick to their guns, rejecting my three plausibility conditions and insisting that Phyllis is obligated at present to administer two doses of Drug C. But then I think that they owe us some account of how deontic principles differ from evaluative principles that's more plausible than the one that I've presented above. We already know that it would be better for Phyllis to administer two doses of Drug C rather than two doses of Drug P; evaluative principles tell us that much. So the question is: what makes the claim that Phyllis is obligated at present to administer two doses of Drug C (and prohibited from administering two doses of Drug P) other than superfluous? That is, what role does this claim play beyond ranking the administration of two doses of Drug C above the administration of two doses of Drug P? I'm at lost as to what that role might be if not the ones that I've specified above.

4. Why PYBSSO is just right

So we've seen that any version of $\text{PYB}\times\text{O}$ that is as broad as, or broader than, PYBPO is too broad and that any version of $\text{PYB}\times\text{O}$ that is as narrow as, or narrower than, PYBSO is too narrow. And we've seen that that PYBSSO falls in the middle, as it is narrower than PYBPO but broader than PYBSO. But, of course, this in itself doesn't establish that PYBSSO is just right. To show that PYBSSO is just right, I must show that anything broader than PYBSSO is too broad and that anything narrower than PYBSSO is too narrow. Below, I'll try to do just that. I'll focus on blameworthiness (and I'll use 'blameworthy' to mean 'worthy of either rational criticism and/or moral blame'), but similar arguments

with respect to predictability could also be made. I'll start with an argument for why anything broader than PYBSSO is too broad:

- 1a If the range of options that an agent can be obligated at t to perform were broader than those that are scrupulously securable by her at t , then an agent who had, at all times leading up to and including time t , all those attitudes that she had decisive reason to have and only those attitudes that she had sufficient reason to have could be blameworthy for failing to fulfill an obligation that she supposedly had at t .
- 2a An agent who had, at all times leading up to and including time t , all those attitudes that she had decisive reason to have and only those attitudes that she had sufficient reason to have could not be blameworthy for failing to fulfill an obligation that she supposedly had at t .
- 3a Therefore, the range of options that an agent can be obligated at t to perform is not broader than those that are scrupulously securable by her at t .³⁴

The argument is deductively valid, and its first premise (viz., 1a) follows directly from SSO. To see this, consider that if S could be obligated at t to ϕ at t' even though her ϕ -ing at t' was not, as of t , scrupulously securable by her, then S could have obligations at t that she would not fulfill even if she "were at t aware of all the relevant reason-constituting facts and were to respond appropriately to them, thereby coming to have at t' all those attitudes that, given those facts, she has decisive reason to have and only those attitudes that she has, given those facts, sufficient reason to have." That is, S could be obligated at t to ϕ at t' even though she would not ϕ at t' even if she were to have at t all those attitudes that she has decisive reason to have and only those attitudes that she has sufficient reason to have. And since the explanation for why she would not ϕ at t' even if she were to have at t all those attitudes that she has decisive reason to have and only those attitudes that she has sufficient reason to have may not lie with any past failure to have the appropriate attitudes, but instead lie with some innate predisposition or to the external manipulations of some nefarious neurophysiologist, it follows that "if the range of options that an agent can be obligated at t to perform were broader than those that are scrupulously securable by her at t , then an agent who had, at all times leading up to and including time t , all those attitudes that she had decisive reason to have and only

³⁴ This is a revised version of an argument that I gave in Portmore 2011a, p. 173.

those attitudes that she had sufficient reason to have could be blameworthy for failing to fulfill an obligation that she supposedly had at *t*.”

The only possible point of contention, then, lies with premise 2a. But there should be no dispute over 2a. For if an agent who had, at all times leading up to and *including* time *t*, all those attitudes that she had decisive reason to have and only those attitudes that she had sufficient reason to have could be blameworthy for failing to fulfill an obligation that she supposedly had at *t*, then being perfectly rational at all times leading up to and *including* time *t* would be insufficient to ensure that she doesn't violate any of the rational obligations that she supposedly has at *t*. And that's implausible. To illustrate, consider a variation on *Temptation* in which we leave everything the same but add the assumption that Phyllis is ideally informed and has been and will continue to be perfectly rational up until Tuesday morning, which is when she will then irrationally decide to spend the day alone rather than meet up with her sponsor. Assume, then, that her alcoholism is not due to any mistake in her past but is instead due to the external manipulation of her brain by the nefarious Dr. White.

So assume that it's just before 1 P.M. Monday and Phyllis must decide which drug to administer to Patience. Now, she has had and continues to have all those attitudes that she has decisive reason to have and only those attitudes that she has sufficient reason to have. Yet nothing she does or intends to do now is going to change the fact that she is going to make the future mistake of choosing to spend tomorrow alone, which will result in her ending up passed out drunk and unable to administer a dose of Drug C on Tuesday.³⁵ Thus, the only way for Phyllis to prevent Patience's death and her own life-long incarceration is to administer Drug P at 1 P.M. Monday. Being perfectly rational, Phyllis predictably takes the only rational course of action and administers Drug P at 1 P.M. Monday. Yet, according to PYBPO, she had an obligation to administer Drug C at 1 P.M. Monday and is, therefore, blameworthy (that is, appropriately criticizable) for having responsibly administered Drug P instead. But how could it be appropriate to criticize her for responsibly having administered Drug P at 1 P.M. Monday when that is surely what any perfectly rational and ideally informed agent would do in her circumstances. Any ideally informed agent would know that administering Drug C would result in death for Patience and life imprisonment for herself. And any perfectly rational agent would prefer a partial

³⁵ I'm assuming that it is possible for Phyllis to have been perfectly rational at all times leading up to and including time *t* and yet be such that she will make some future mistake regardless of what she does or intends to do at *t*. If the reader finds this possibility hard to imagine, then just assume that the nefarious Dr. White has manipulated her brain such that she is causally determined to make this future mistake.

cure for Patience and loss of one's job (but not one's freedom) to that outcome. So I think that we must accept that the above argument is sound and conclude that anything broader than PYBSSO is too broad.

We should also hold, I believe, that anything narrower than PYBSSO is too narrow. Here's my argument:

- 1b** If the range of options that an agent can be obligated at t to perform were narrower than those that are scrupulously securable by her at t , then an agent who ϕ -ed when it would have been better to have ψ -ed could escape blame/criticism for having failed to have ψ -ed merely because she lacked the background attitudes that she had decisive reason to have and that, if possessed by her, would have made her able to intend effectively to ψ .
- 2b** An agent who ϕ -ed when it would have been better to have ψ -ed cannot escape blame/criticism for having failed to have ψ -ed merely because she lacked the background attitudes that she had decisive reason to have and that, if possessed by her, would have made her able to intend effectively to ψ .
- 3b** Therefore, the range of options that an agent can be obligated at t to perform is not narrower than those that are scrupulously securable by her at t .

This argument is also deductively valid, and its first premise (viz., 1b) also follows directly from SSO. The only possible point of contention, then, lies with 2b. But we should accept 2b. The idea behind 2b is that not being able to intend effectively to ψ because one continues to maintain an inappropriate set of background attitudes is not an adequate excuse for failing to ψ . In *Foolishness*, for instance, Phyllis cannot excuse herself from having to administer two doses of Drug C simply because she cannot, given her current background attitudes, effectively intend to administer the second dose of Drug C. For all she needs to do is to respond appropriately to her reasons and thereby form the intention to meet up with her sponsor tomorrow, as she ought to. And with this attitude in the background, she can then effectively intend to administer the second dose of Drug C.

So we should accept both 3a and 3b, and that means that we should think that PYBSSO is just right.

5. Potentially Problematic Cases for PYBSSO

Below, I consider three potentially problematic cases for PYBSSO and argue that none of them pose a serious challenge to PYBSSO.

5.1 How the addition of new information can affect one's reasons

Consider the following case:

The Two Boxes: Before Jones are two boxes: A and B. He may choose at most one of the two, and he will then receive however much money is in the box that he chooses. Jones knows that there's only \$1 in B but that there's a lot of money in A. Indeed, he knows that there is more than \$1,000,000 in A. But he doesn't know exactly how much money is in A. There is in fact \$1,786,444.23 in A. The nefarious Dr. White, however, is on the scene and will detonate a bomb killing Jones if he chooses A while at the same time knowing that there is \$1,786,444.23 in A.³⁶

The fact that there is \$1,786,444.23 in A constitutes an objective reason for Jones to believe that there is \$1,786,444.23 in A. Thus, if Jones were aware of all the reason-constituting facts and were to respond appropriately to them, he would know that there is exactly \$1,786,444.23 in A. And, knowing this, his best option would be to choose B. For it is better to receive \$1 and live than to receive \$1,786,444.23 and die. So if we were to define 'scrupulously securable' in terms of the awareness of *all* the reason-constituting facts, it would turn out that Jones's choosing A would not be Jones's best scrupulously securable option. Indeed, it would be his worst, for his only scrupulous securable options would be: (1) choose B and go on living with an extra \$1, (2) choose neither and go living with no extra money, and (3) choose A and receive \$1,786,444.23 but die before having the chance to spend any of it. Thus, if we were to combine PYBSSO with a definition of 'scrupulously securable' in terms of full information, we would get the counterintuitive result that Jones is obligated, in his current circumstances of not knowing how much exactly is in A, to choose B. Yet, given his current epistemic situation, Jones has the opportunity to go on living with an extra

³⁶ I borrow this example from Peter Graham (correspondence 2011). He presented this counterexample when I had initially proposed a subjunctive analysis of 'scrupulously securable' in terms of full information and perfect rationality. Jacob Ross (*forthcoming*) presents a similar case called the *Ice Cream of Doom*. It too is a case in which the agent's coming to have full information affects what she has reason to do. I focus on *The Two Boxes*, because this example is simpler. But my response to the *Ice Cream of Doom* is the same as my response to *The Two Boxes* below.

\$1,786,444.23 to spend by choosing A. And the fact that his better informed counterpart would die if he were to choose A seems irrelevant when it comes to determining what the actual Jones should do given his actual circumstances.

It's fortunate, then, that I avoid this problem in that I define 'scrupulously securable', not in terms of the agent's awareness of *all* the reason-constituting facts, but in terms of the agent's awareness of the *relevant* reason-constituting facts—see SSO above. What are the relevant reason-constituting facts? Well, let S^* be S 's fully informed counterpart. Thus, S^* knows in addition to whatever else S knows all other truths.³⁷ Now, for any true proposition p , the fact that p is a relevant reason-constituting fact with regard to S 's ϕ -ing at t'' if and only if a perfectly rational S^* would want S to know that p when deliberating about whether or not to ϕ at t'' .³⁸ Thus, the fact that there is exactly \$1,786,444.23 in A would not count as a relevant reason-constituting fact, for a perfectly rational Jones* would not want the actual Jones to know that there is \$1,786,444.23 in A given that this would preclude the actual Jones from being able to gain an extra \$1,786,444.23 to spend. And since the fact that there is \$1,786,444.23 in A is not a relevant reason-constituting fact, it is not the case that Jones would come to have the belief that there is \$1,786,444.23 in A if Jones were at t aware of all the *relevant* reason-constituting facts and were to respond appropriately to them. Therefore, according to SSO, Jones's choosing A at t'' is his best scrupulously securable option at t . And, thus, PYBSSO rightly implies that Jones's is obligated to choose A at t'' .

5.2 The permissibility of acting so as to avoid future obligations to do good

Consider the following case:

Selfish Sid: Sid has exactly \$500. Unless Sid uses this \$500 to purchase something for himself today, he will face the following choice tomorrow: use the \$500 on that day either to purchase something for himself or to save three innocent lives. Assume that Sid's using the \$500 tomorrow to save the three innocent lives is not as of today scrupulously securable by him. For, if he still has the \$500 tomorrow, he will, given his selfishness, succumb to the temptation to use it to purchase the latest and greatest tablet computer for himself when it debuts tomorrow, and this is true regardless of what his

³⁷ S^* is different than what I referred to as ' S^+ ' above. Whereas S^* knows all the facts, S^+ knows only all the *relevant* reason-constituting facts.

³⁸ This is modeled after Michael Smith's advice model of the internalism requirement—see M. Smith 1995.

actions and intentions are today. Thus, even if he were to form the present intention (or resolution) to use the money tomorrow to save the three lives, temptation would get the best of him tomorrow, and he would then end up purchasing the tablet computer, letting the three die. Nevertheless, his not spending the money today is scrupulously securable by him at present. All he has to do is to intend now to spend the \$500 on the tablet computer when it debuts tomorrow. And, of course, once tomorrow rolls around, his using the \$500 (if he still has it) to save the three lives will be scrupulously securable by him at that point. For, at that point, all he has to do to secure the survival of the three is to intend to save them. It's just that that's not what he's going to intend to do if he is put in that situation. Thus, unless he spends the money today, he'll end up doing something wrong tomorrow. Knowing this, Sid spends the money today on some other cool electronic gadget. (Assume that what Sid has most reason to do, all things considered, is to perform his morally best option.)³⁹

Unless we are to suppose that Sid has better reason to act so as to put himself in the position of having a future obligation that he won't fulfill than to act so as to avoid incurring that future obligation, PYBSSO implies that Sid is permitted to spend the \$500 on himself today, thereby avoiding an obligation to save the three tomorrow. Some, such as Timmerman (2011), find this problematic, but I don't see why. That is, I don't see why we should think that it would be wrong for Sid to spend the \$500 on himself today unless we think that he has better reason to act so as to put himself in the position of having a future obligation that he won't fulfill.

In any case, I think that we can see why we should reject Timmerman's judgment that Sid ought to refrain from using the \$500 today by considering a slight variant on *Selfish Sid*. So imagine *Selfish Sid 2*, which is just like the original except that in this case Sid has *today* the option of using the \$500 either to purchase something for himself or to save *two* innocent lives. The options tomorrow remain the same. So, in *Selfish Sid 2*, he has on both days the choice of using the \$500 for his own benefit or using it to save lives. It's just that he has the opportunity to save only two lives today but will have the opportunity to save three tomorrow (that is, if he doesn't use the \$500 today). Given the stipulation that Sid won't save the three even if he were put in position to do so tomorrow, it seems clear that Sid should take this opportunity to save the two. There's no point in forgoing the opportunity to save the two today just so that he can be in a position to save the three tomorrow, when, as a matter of fact, Sid won't save the

³⁹ This example is adopted from Timmerman 2011.

three tomorrow even if he is put in that position. It seems better for Sid to use the \$500 today to save the two. So I don't see a problem here for PYBSSO. We should accept, as PYBSSO implies, that Sid has no obligation to refrain from using the money today so as to put himself in the position of doing some (or more) good tomorrow, when in fact he won't do that (greater) good and there's nothing that he can do now to change that.⁴⁰

5.3 Presently obligatory future wrong-doing

Even if I'm right in thinking that there is nothing problematic about PYBSSO's implying that it is permissible to act so as to avoid future obligations, PYBSSO has some other potentially more troubling implications. PYBSSO implies that agents are sometimes presently obligated to do something in the future that they will be obligated to refrain from doing when that future time rolls around. To illustrate, consider the following case, which I borrow from Ross (*forthcoming*):

"Satan's School for Girls: It is now June 6, 2011. Sally is a fine, upstanding teenage girl whose moral character is impeccable. However, she is about to be kidnapped by Satanists, and brought up in Satan's School for Girls. As a result, her moral character will be severely corrupted, and she will come to desire, more than anything else, to sacrifice her firstborn child to the Prince of Darkness. At the appointed hour, on June 6, 2016, Sally will have the opportunity to kill her firstborn child, and will be able to do so either by cutting off the child's head with an axe, or by bludgeoning the child to death with a club. At that very moment, Child Services personnel will be arriving on the scene, and so if she refrains from killing her child, the child will be taken into protective custody and live happily ever after. If, on June 6, 2011, before being kidnapped by the Satanists, Sally were to intend to cut off the head of her firstborn child, then she would follow through with this intention five years later. And if she were to intend to bludgeon her child to death, then she would likewise follow through with this intention five years later. But if she were to intend not to kill her firstborn child, then she would change her mind after becoming a Satanist, and end up cutting off her child's head. The above three subjunctive conditionals are true not by any fault of Sally's, but purely in virtue of the fact that she is about to be kidnapped and brought up at Satan's School for Girls." Assume that, "among the two ways of killing her firstborn child, cutting off its head would cause less pain."

⁴⁰ Nor, I'm assuming, is there any set of attitudes that he has sufficient reason to have that would change this.

PYBSSO implies that Sally is obligated at present (on June 6, 2011) to cut off the head of her firstborn child on June 6, 2016, and yet PYBSSO also implies that, no matter what she may do in the intervening five years, she will be obligated on June 6, 2016, to refrain from cutting off the head of her firstborn child on that day. Thus, PYBSSO implies that she is obligated at present to do something on June 6, 2016, that come June 6, 2016, she will be obligated as of at that time to refrain from doing. Ross calls this the “problem of nonratifiability,” because it involves making a *nonratifiable prescription*: that is, PYBSSO makes a prescription prior to the time of action that it “will inevitably reverse at the time of action, regardless of what the agent does in the interim.” Ross finds this to be “a very peculiar implication.”

The only way to avoid this implication while accepting some version of PYBxO is to either (A) accept some version of PYBxO that implies that Sally has *at present* the option to refrain from killing her child five years from now despite the fact that she will do so regardless of what her present actions and attitudes are or (B) accept some version of PYBxO that implies that Sally won’t have the option five years from now of refraining from killing her child despite its being true that she would refrain from killing her child five years from now if she were at that time to form the intention to refrain from killing her child—an intention that she will at that time have the capacity to form. Both are unacceptable.

First, consider option B. Assume that the “appointed hour” is noon. Since Ross stipulates that Sally’s refraining from killing her child at noon on June 6, 2016, will be a *scrupulously securable option* for Sally at 11:59 A.M. on June 6, 2016, it follows both that (a) Sally would refrain from killing her child at noon on June 6, 2016, if she were to intend at 11:59 A.M. on June 6, 2016, to refrain from killing her child at noon on June 6, 2016, and that (b) Sally has at 11:59 A.M. on June 6, 2016, the capacity to continue, or to come, to have at noon on June 6, 2016, the intention to refrain from killing her child at noon on June 6, 2016. This means that Sally’s refraining from killing her child at noon on June 6, 2016, is an *x* option for Sally at 11:59 A.M. on June 6, 2016, on every version of PYBxO that we’ve considered above. Indeed, I can think of no plausible version of PYBxO in which this isn’t an *x* option for Sally at 11:59 A.M. on June 6, 2016. So we should reject option B, leaving us with only option A.

So consider option A. First, it seems clear, given the arguments above, that no version of PYBxO that entails that Sally has *at present* the obligation (and, thus, the option) to refrain from killing her child five years from now is going to meet any of the three plausibility conditions. Take the practical condition, for instance. There is no way that Sally could use any such version of PYBxO at present as a guide to fulfilling the present obligation that she putatively has to refrain from

killing her child five years from now. After all, Ross stipulates that Sally will kill her child five years from now regardless of her present actions, intentions, and other attitudes. Of course, Ross just finds it absurd to think that Sally is obligated at present to do something on June 6, 2016, that come June 6, 2016, she will be obligated as of that time to refrain from doing. And my rehashing the arguments in favor of PYBSSO that appeal to the three plausibility may do nothing to change that fact. But perhaps I can shake his confidence by appealing to the following revised version of the example: *Satan's School for Girls II*. Imagine that in this case the brainwashing will result only in her having a desire to harm her child, but not necessarily a desire to kill her child. And assume that if she intends now to prick her daughter's finger with a pin than she will later on follow through with only that intention. But that if she intends now not to harm her child, she will after being brainwashed change her mind and decide to bludgeon her child to death. In this case, it doesn't seem absurd at all to suppose that at present Sally is obligated to prick her child's finger five years from now but that come June 6, 2016, she should refrain from harming her child in any way given that at that point she can secure that by merely responding appropriately to her reasons.

Hopefully, this example will shake Ross's confidence in the claim that non-ratifiable prescriptions are absurd. It does mine. In any case, any theory (such as PYBSSO) that holds that one's options *at a time* can broaden as well as narrow over time will imply that there are non-ratifiable prescriptions.⁴¹ And I don't see any reason to think that options at a time couldn't broaden as well as narrow.⁴² Of course, there are versions of PYBxO (such as PYBPO) where options at a time can only narrow. But all such versions that we've considered have turned out to be unacceptable. So I'm not convinced that *Satan's School for Girls* poses a serious challenge to PYBSSO. Indeed, I'm inclined to think that options at a time can broaden as well as narrow over time and that, therefore, there can be non-ratifiable prescriptions. And absent Ross or someone else proposing some

⁴¹ No one denies that there is a certain sense in which one's options can broaden over time. That is, no one denies the following is possible: earlier at t_1 there were only two things that I could do (at t_1), but now at t_2 , after my bonds have been broken, there are now three things that I can do (at t_2). Some do, however, deny that one's options *at a time* can broaden over time. That is, they deny that the following is possible: at t_1 there were only two things that I could do at t_3 , but at t_2 there are now three things that I could do at t_3 . Yet these same people concede that one's options *at a time* can narrow over time. That is, they concede that the following is possible: at t_1 there were three things that I could do at t_3 , but at t_2 there are now only two things that I could do at t_3 . Thanks to Peter Graham for pointing out the need to clarify this.

⁴² Goldman (1978, p. 196) also holds that options at a time can broaden as well as narrow.

version of PYBxO that avoids non-ratifiable prescriptions while also meeting the above three plausibility conditions, I'm inclined to stick with my conclusion that PYBSSO is the most plausible version of PYBxO on offer.

6. Conclusion

I've argued that SSO is the Goldilocks of accounts of our options. Anything broader is too broad. Anything narrower is too narrow. And, so, SSO is just right. This means that insofar as we should accept PYBxO, we should accept PYBSSO. My argument is based on the assumption that there are certain conditions on the plausibility of any proposed deontic principle. It is open to those who want to reject PYBSSO to deny these conditions, but if they do, they owe us some account of what role deontic principles play over and above the role of ranking actions from better to worse, a role that is already more than adequately played by evaluative principles. If deontic principles are just another (more coarse-grained) way of ranking some actions as better (or worse) than others, then, given that we already have evaluative principles that provide a maximally fine-grained way of ranking actions from better to worse, deontic principles turn out to be entirely superfluous. And given that we do have a need to be able to assess whether agents are criticizable and to predict how agents may act insofar as they are both ideally informed and perfectly rational and that deontic principles seem well-suited to this task, I think that we should be loath to condemn deontic principles to being superfluous short of there being some powerful reason for thinking that we have no choice but to consider them superfluous.⁴³

References

- Bergström, L. (1977). "Utilitarianism and Future Mistakes," *Theoria* 43: 84–102.
- Bratman, M. E. (2009). "Intention, Belief, Practical, Theoretical." In S. Robertson (ed.), *Spheres of Reason*, pp. 29–62. Oxford: Oxford University Press.
- Chappell, R. Y. (*forthcoming*). "Fittingness: The Sole Normative Primitive," *Philosophical Quarterly*.

⁴³ For help with this paper, I thank Dale Dorsey, Peter Graham, Brian Hedden, and Travis Timmerman.

- Dancy, J. (2004). "Enticing 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. 91–118. Oxford: Oxford University Press.
- Feldman, F. (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht: D. Reidel Publishing Company.
- . (1975). "World Utilitarianism," In K. Lehrer (ed.), *Analysis and Metaphysics*, pp. 255–271. Dordrecht: D. Reidel Publishing Company.
- Fischer, J. M., and Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Gert, J. (2004). *Brute Rationality: Normativity and Human Action*. Cambridge: Cambridge University Press.
- . (2003). "Requiring and Justifying: Two Dimensions of Normative Strength." *Erkenntnis* 59: 5–36.
- 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.
- Hare, C. (2011). "Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had not Done What You Did," *Noûs* 45: 190–206.
- Harman, G. (1976). "Practical Reasoning," *Review of Metaphysics* 29: 431–463.
- Hedden, B. (2011). "Options and the Subjective Ought." Unpublished manuscript, Massachusetts Institute of Technology, Cambridge, MA.
- Jackson, F. (1985). "On the Semantics and Logic of Obligation," *Mind* 94: 177–195.
- Jackson, F. and Pargetter, R. (1986). "Oughts, Options, and Actualism," *The Philosophical Review* 95: 233–255.
- Mele, A. (1989). "She Intends to Try," *Philosophical Studies* 55: 101–106.
- Norcross, A. (2006). "Reasons without Demands: Rethinking Rightness." In J. Dreier (ed.), *Contemporary Debates in Moral Theory*, pp. 38–53. Oxford: Blackwell.
- Portmore, D. W. (2011a). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.
- . (2011b). "Perform Your Best Option." Unpublished manuscript, Arizona State University, Tempe, AZ.

- . (2011c). "The Teleological Conception of Practical Reasons," *Mind* 120: 117–153.
- Raz, J. (1999). *Engaging Reason: On the Theory of Value and Action*. Oxford: Oxford University Press.
- Ross, J. (forthcoming). "Actualism, Possibilism, and Beyond." *Oxford Studies in Normative Ethics*.
- . (2009). "How To Be a Cognitivist about Practical Reason." *Oxford Studies in Metaethics* 4: 243–282.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, Mass.: Belknap Press.
- Smith, H. M. (1988). "Making Moral Decisions," *Noûs* 22: 89–108.
- Smith, M. (1995). "Internal Reasons," *Philosophy & Phenomenological Research* 55: 109–131.
- Sobel, J. H. (1982). "Utilitarian Principles for Imperfect Agents," *Theoria* 48: 113–126.
- . (1976). "Utilitarianism and Past and Future Mistakes," *Noûs* 10: 195–219.
- Timmerman, T. (2011). "Why There Is a Relevant Difference between Inter- and Intra-Agent Act-Sets in Terms of Moral Responsibility." Unpublished manuscript, Syracuse University, Syracuse, NY.
- Velleman, J. D. (1989). *Practical Reflection*. Princeton: Princeton University Press.
- Vessel, J.-P. (2003). "Counterfactuals for Consequentialists," 112: 103–125.
- Woodard, C. (2008). *Reasons, Patterns, and Cooperation*. New York: Routledge.