

Moral Obligations: Actualist, Possibilist, or Hybridist?

Travis Timmerman and Yishai Cohen
Published in the *Australasian Journal of Philosophy*

Do facts about what an agent would freely do in certain circumstances at least partly determine any of her moral obligations? Actualists answer ‘yes’ while possibilists answer ‘no’. We defend two novel hybrid accounts that are alternatives to actualism and possibilism: *Dual Obligations Hybridism* and *Single Obligation Hybridism*. By positing two moral ‘oughts’, each account retains the benefits of actualism and possibilism, and yet is immune from the *prima facie* problems that face actualism and possibilism. We conclude by highlighting one substantive difference between our two hybrid accounts.

Keywords hybridism; actualism; possibilism; securitism; obligations; intentional control

1. Introduction

Actualists hold that what an agent would freely do in certain circumstances at least partly determines her moral obligations. Possibilists disagree. To illustrate these views in further detail, consider the following case adapted from Michael Zimmerman [2006: 153]:

The Wedding Invitation

Your ex-partner invites you to her wedding, and it would be best if you accepted the wedding invitation and attended without misbehaving—something you *can* do at little cost to yourself. But due to your jealousy, if you were to accept the wedding invitation, you would misbehave at the wedding. Thus, it would be better to decline the invitation rather than to accept it and misbehave. But declining would be worse than accepting and behaving.

Are you morally obligated to accept the invitation? According to actualists, because what *would* actually happen if you were to accept the invitation is worse than what *would* actually happen if you were to decline, you are obligated to decline. Actualism may thus be understood as the view that an agent ought to ϕ just in case what *would* happen if she were to ϕ is better than what *would* happen if she were not to ϕ [Goldman 1976; Sobel 1976; Jackson and Pargetter 1986]. According to possibilism, you ought to accept the invitation. This is because, after accepting the invitation, you *can* (it’s *possible* for you to) attend the wedding and behave yourself. That’s just not what you *would* do if you were to accept the invitation. Possibilism may thus be understood as the view that an agent ought to ϕ just in

case ϕ -ing is part of the best act-set she *can* perform [Greenspan 1978; Feldman 1986; Zimmerman 1996: ch. 6].

While there are a variety of technical formulations of both actualism and possibilism on offer, we will settle with the following:¹

Actualism At t an agent S is obligated to ϕ at t' iff S can ϕ at t' , and what would happen if S were to ϕ at t' is better than what would happen if S were to $\sim\phi$ at t' .

Possibilism At t an agent S is obligated to ϕ at t' iff ϕ -ing at t' is part of the best² maximally specific act-set³ that S can perform.

Both actualism and possibilism are subject to certain *prima facie* problems to be discussed below. In this paper, we offer a novel solution to the actualism/possibilism debate that avoids said problems. Specifically, we argue for a hybrid view that posits two moral ‘oughts’ that jointly track the insights of both actualism and possibilism. Rather than arguing for a particular hybrid view, we canvass two promising views and argue that the solution is somewhere in this vicinity.

Our paper is structured as follows. In the next section we highlight actualism’s and possibilism’s *prima facie* problems. In section three, we consider a preliminary hybrid position that takes into account *both* actualist and possibilist intuitions. In section four, we explain why even this hybrid position requires further crucial modifications, which leads to section five where we articulate and defend *Dual Obligations Hybridism (DOH)*, according to which there are two moral obligations.

¹ For simplicity’s sake, we ignore complications that arise from one’s ability to perform supererogatory acts. Moreover, we are first and foremost concerned with objective obligations rather than subjective ones (as is standard in the actualist/possibilist literature). Roughly, an objective obligation is what an agent should do if she were aware of all of the normatively relevant facts. A subjective obligation, by contrast, is what a conscientious moral agent should do relative to the evidence in her epistemic ken. [cf. Zimmerman 1996: 10–20; Portmore 2011: 10–22].

² If more than one maximally specific act-set is tied for the best, then S is obligated to ϕ iff ϕ -ing is a member of all of the maximally specific act-sets that are tied for the best.

³ A maximally specific act-set that an agent can perform is one that extends from the present moment to the last moment of time at which an agent can perform an action. Note, moreover, that all subsequent discussion of an act-set may refer to a singleton act-set.

In section six, we articulate and defend *Single Obligation Hybridism (SOH)*, according to which there are two moral ‘oughts’, but only one of them is a moral obligation. We end the paper by highlighting one substantive difference between *DOH* and *SOH*.

2. Problems for Actualism and Possibilism

Two important *prima facie* problems with actualism, and one with possibilism, have been identified in the literature. To illustrate these *prima facie* problems, we present another actualist/possibilist scenario that will guide us throughout the rest of this paper.⁴

Evil Professor Smith is asked to review a graduate student’s paper. Smith can perform the following three act-sets:

- X.** (a) agree to review the paper, and (b) review the paper.
- Y.** (~a) decline to review the paper, (~b) not review the paper, (c) plagiarize the paper, and (d) convince the student to quit philosophy to avoid being caught plagiarizing.
- Z.** (a) agree to review the paper, (~b) not review the paper, (c) plagiarize the paper, and (e) murder the graduate student to avoid being caught plagiarizing.

Of act-sets *X*, *Y*, and *Z*, **X** has the **best** possible outcome, viz. the graduate student’s paper getting published in a top-tier journal; **Y** has the **second best** possible outcome; **Z** has the **worst** possible outcome. While Smith can *X*, *Y*, or *Z*, given Smith’s rotten moral character, the following two counterfactuals are true:

1. If Smith were to (a), she would *Z*.
2. If Smith were to (~a), she would *Y*.

Actualists hold that Smith is obligated to (~a), whereas possibilists hold that Smith is obligated to (a).

Now, the first *prima facie* problem with actualism is that it permits an agent *S* to avoid incurring *any*

⁴ This case is loosely adapted from Goldman [1978], variations of which have appeared in the literature, such as in Jackson and Pargetter [1986: 235], Vorobej [2000: 131–2], Portmore [2011: 180], and Timmerman [2015: 1512].

moral obligation to ϕ , which S can easily fulfill, simply in virtue of S 's rotten moral character [Jackson and Pargetter 1986: 240; Zimmerman 2006: 156; Portmore 2011: 207; Baker 2012: 642–43; Timmerman 2015]. To illustrate, actualism holds that Smith is obligated to $(\sim a)$ since, given (1)–(2), (a) -ing would result in the worst possible outcome, viz. Z which includes (e) murdering the student to avoid being caught plagiarizing. However, Smith is intuitively obligated to (a) since she is intuitively obligated to X , and X -ing requires (a) -ing.⁵

Second, actualism prescribes bad behavior, and acting on such prescriptions presumably renders an agent S immune from moral criticism, even when S can easily refrain from such behavior [Wedgwood 2009; Ross 2013].⁶ To illustrate, given (1)–(2), and given that Smith can either (a) or $(\sim a)$, Smith would either Y or Z (though Smith can X , Y or Z). Since Y is the better option of the two, actualism holds that, *even though* Smith can X , Smith is obligated to $(\sim a)$ which is, we may suppose, an instance of bad behavior.

Third, possibilism entails that an agent S has an action-guiding obligation to ϕ even if S would perform a subsequent act-set that is deeply morally wrong, and worse than the act-set S would subsequently perform if S were to $\sim\phi$ [Goldman 1976: 469–70; Woodward 2009: 219–20; Portmore 2011: 211]. To illustrate, possibilism holds that Smith is obligated to X , and thus is obligated to (a) , since X is the best act-set that Smith can perform. Moreover, Possibilism prescribes (a) -ing, *even though* given (1), it follows that if Smith were to (a) , she would Z , which includes (e) murdering the grad student to avoid being caught plagiarizing. With an understanding of these three *prima facie* problems in place, we are now in a position to develop a more satisfactory view of the nature of moral obligations.

⁵ See Portmore [2011: 181–3; 2013] for a defense of the ‘permissibility distributes over conjunction’ principle implicitly invoked here.

⁶ Although most actualists seem to make this assumption, Goldman’s [1978: 207] revised version of actualism might be the sole exception. Nevertheless, Goldman’s view is still subject to the first problem we identified with actualism. Moreover, Goldman’s view has the following implausible implication: an agent who is aware of all of the normatively relevant facts who fulfills all of her moral obligations for the right reasons can still be morally criticizable.

3. A First Pass at a Hybrid View

Scenarios like *The Wedding Invitation* and *Evil Professor Smith* seem to generate conflicting intuitions that pull us towards both actualism and possibilism. For, consider the central tension that is highlighted in *Evil Professor Smith*. Given (1)–(2), irrespective of whether Smith performs (a) or ($\sim a$), Smith seemingly acts in a morally dubious manner. On the one hand, since Smith can X , Smith seems obligated to X (and X -ing requires (a)-ing). But on the other hand, since Smith would murder the grad student if she were to (a), then Smith seems obligated to ($\sim a$). So, there is certainly an intuitive pull towards *both* actualism and possibilism. But both views cannot be true. Michael McKinsey [1979: 388] goes so far as to say that the actualist/possibilist scenarios generate a paradox.

In light of these apparently conflicting intuitions, we think that only a hybrid view that posits *two* moral ‘oughts’ can show why the paradox of actualist/possibilist scenarios is illusory. So, according to a preliminary proposal, there are two moral obligations—one actualist and one possibilist. Of course, one will immediately wonder which obligation is action-guiding since there are many instances in which one cannot act on *both* obligations. A natural suggestion is that the actualist obligation is action-guiding since actualism’s prescriptions take into account facts about an agent’s moral shortcomings (we will expound upon this action-guiding notion in the next section). In other words, only an actualist obligation takes into account facts about how agents would subsequently act and consequently avoids the problem that would result from making the possibilist obligation the action-guiding one, viz. requiring an agent to do something, even when she would perform a subsequent act that is deeply morally wrong.

On the view just proposed, agents have both an actualist and a possibilist obligation. Additionally, when one’s actualist and possibilist obligations conflict, one’s actualist obligation is action-guiding. Let’s call this view *Simple Hybridism*. In regards to *Evil Professor Smith*, *Simple Hybridism* holds that although Smith has a possibilist obligation to X and thus has a possibilist obligation to (a), Smith also has an action-guiding actualist obligation to ($\sim a$), precisely because of

(1)–(2). *Simple Hybridism* is not only able to account for the fact that, in one sense, Smith is obligated to (a) , and in another sense, Smith is obligated to $(\sim a)$; it's also immune from the three *prima facie* problems that face actualism and possibilism, as we will now demonstrate.

The first problem we noted with actualism is that it permits an agent S to avoid incurring *any* moral obligation to ϕ , which S can easily fulfill simply in virtue of having a rotten moral character. Actualism holds that Smith is morally obligated to $(\sim a)$ since, given (1)–(2), (a) -ing would result in the worst possible outcome, viz. Z . By contrast, according to *Simple Hybridism*, although Smith does avoid incurring an actualist obligation to (a) , Smith still has a possibilist obligation to (a) . So, while *Simple Hybridism* permits an agent S to avoid incurring *one* moral obligation to (a) in virtue of having a rotten moral character, it's not true on *Simple Hybridism* that S can avoid incurring *any* moral obligation to (a) in virtue of having a rotten moral character.

The other problem with actualism is that it prescribes bad behavior, and acting on such prescriptions presumably renders S immune from moral criticism, even when S can easily refrain from such behavior. Given (1)–(2), and given that Smith can either (a) or $(\sim a)$, Smith would either Y or Z . Since Y is the better option of the two, actualism holds that Smith is obligated to $(\sim a)$, *even though* Smith can easily X . By contrast, *Simple Hybridism* does not prescribe bad behavior for Smith, *such that acting on this prescription renders Smith immune from moral criticism, even when Smith can easily refrain from such behavior*. To be sure, *Simple Hybridism* does hold that Smith has an action-guiding, actualist obligation to $(\sim a)$. But Smith still has a possibilist obligation to (a) and could be criticizable for $(\sim a)$ -ing.

Lastly, according to possibilism, Smith is obligated to X , and thus is obligated to (a) , *even though* given (1), it follows that if Smith were to (a) , she would Z , which includes (e) murdering the grad student to avoid being caught plagiarizing. By contrast, *Simple Hybridism* *does* take into account the fact that Smith would do something deeply wrong if she were to (a) . This is precisely why Smith has an action-guiding, actualist obligation to $(\sim a)$.

4. Problems with *Simple Hybridism*

Simple Hybridism has been shown to accommodate the seemingly conflicting intuitions that arise in actualist/possibilist scenarios, while simultaneously escaping the three *prima facie* problems that face actualism and possibilism. Nevertheless, *Simple Hybridism* faces two of its own problems.

The first problem with *Simple Hybridism* is that it may appear *ad hoc*. The view's two obligations appear unrelated to one another, despite the fact that they are both *obligations*. As a result, it would be better to have a view that retains all of *Simple Hybridism*'s advantages, but doesn't posit two seemingly unrelated types of moral obligations.

Here is the second problem. The careful reader will notice that actualism prescribes conflicting obligations. Consequently, *Simple Hybridism*'s actualist obligation likewise prescribes conflicting obligations. To see this, recall that according to the aforementioned definition of actualism, there are no restrictions on which act-sets can stand in for ' ϕ ' or what times can stand in for ' t '. As such, although actualism holds that Smith is obligated to $(\sim a)$, actualism *also* holds that Smith is obligated to X since what would happen if Smith were to X is better than what would happen if Smith were not to X . So, given that X -ing requires (a) -ing, it follows that actualism holds that Smith is obligated to $(\sim a)$ *and* that Smith is obligated to X , which requires (a) -ing. Moreover, actualism doesn't say which obligation takes priority [Carlson 1999: 261]. Consequently, the actualist obligation that *Simple Hybridism* adopts faces the same problem.

In order to escape this problem, we should restrict what can stand in for ' ϕ ' to actions that are under an agent's control of some sort. For instance, consider the following notion of control:

ϕ -ing at t' is under an agent S 's *intentional control* at t iff S can⁷ form at t a certain set of intentions, such that if S were to form such intentions at t , S would ϕ at t' .⁸

⁷ Notice that we are not employing a conditional analysis of abilities. Rather, we intend to be officially neutral with respect to the proper analysis of 'can'. We are, however, inclined to hold that if S can perform an act-set, then the following is true: the first basic act of the act-set either is or will be under S 's intentional control and, if

Given this notion of control, we can modify *Simple Hybridism*'s actualist obligation as follows:

At t an agent S has an **actualist* obligation** to ϕ at t' iff ϕ -ing at t' is an act-set under S 's intentional control at t , and what *would* happen if S were to ϕ at t' is better than what would happen if S were to perform any other incompatible act-set under S 's intentional control at t .

We may call a hybrid view that posits both an action-guiding actualist* and (a non-action-guiding) possibilist obligation *Simple Hybridism**. To be clear, a view that posits only an actualist* obligation—call this view actualism*—doesn't escape the two *prima facie* problems to which actualism succumbs. We're now going to provide one more case that illustrates this point and also illuminates the somewhat non-standard action-guiding notion to which we are appealing:⁹

Gamer Gabi

Gabi has exactly \$500 in her account and had planned to use it to purchase a PlayStation 4 (PS4) tomorrow. Unless Gabi spends the \$500 on herself today, she will face the following choice tomorrow: use the \$500 to save three innocent lives or purchase a PS4. It's not under Gabi's present intentional control to both keep the money in her account today *and* use the money to save the three lives tomorrow. However, it *is* under her present intentional control to keep the money in her account today. Moreover, if she does keep the money in her account today it *would be* under her intentional control tomorrow to use the money to save the three

S were to perform that basic act, then the next basic act either is or will be under S 's intentional control, and so on for each of the basic acts that make-up that act-set [cf. Goldman 1978: 197].

⁸ While we think that intentional control is the kind of control that has a bearing upon what an agent morally ought to do in the actualist sense, one could alternatively adopt our hybrid solution and accept a different kind of control, such as Portmore's [2011: 151, 167–9] notion of present deliberative control, or Ross's [2013: 82] notion of a directly securable option. So, here and throughout the paper, one may substitute "intentional control" with some relevantly similar notion of control. To emphasize, our defense of *DOH* and *SOH* doesn't hinge upon adopting our preferred notion of intentional control.

⁹ The *Gamer Gabi* case is a variant of Timmerman's [2015: 1516–7] *Selfish Sid* case. Timmerman argues that (like actualism*) Portmore's Moral Securitism succumbs to actualism's two *prima facie* problems. Interestingly, the same charge could be made against Smith and Ross's positions, since the kind of control one employs in order to restrict what can stand in for ' ϕ ' cannot avoid this problem. That is, any view that posits a single actualist-style obligation that incorporates a kind of control to pick out the relevant act-set that can stand in for ' ϕ ' will render the intuitively incorrect verdict in *Selfish Sid*-style cases. Only positing two moral 'oughts' allows one to avoid said problems.

lives. At that point, all she has to do is intend to use the \$500 to save the three lives. That's just not what she would do if she finds herself in that situation. Finally, Gabi is aware of these facts and consequently decides to purchase a PS4 for herself today rather than tomorrow. She does so with the intention of avoiding any wrongdoing.

Actualism* renders the verdict that Gabi lacks an obligation to save the three lives tomorrow. More specifically, it entails that purchasing a PS4 tomorrow (rather than saving the three lives tomorrow) is wrong, but purchasing a PS4 today (rather than saving the money to save the three lives tomorrow) is permissible. But this seems to get Gabi off the hook too easily. After all, Gabi can secure the saving of the three lives tomorrow by intending today not to spend the money, and then intending tomorrow to save them. Hence, in order to escape the first problem that actualism faces, we need to posit a possibilist obligation in addition to an actualist* obligation, such that Gabi still has an obligation to save the three lives tomorrow. Thus, *Simple Hybridism** renders the intuitively correct verdict that Gabi has an obligation to save the three lives tomorrow.

Actualism* also succumbs to actualism's second *prima facie* problem. To see this, suppose that if Gabi purchases a PS4 today (rather than tomorrow), the deaths of the three people would be painless, whereas if she purchases a PS4 tomorrow, their deaths would be extremely painful. According to actualism*, then, Gabi has an obligation to purchase a PS4 today since, of all the act-sets under her present intentional control, performing that act-set would result in the best outcome. Thus, actualism* prescribes the bad behavior of purchasing a PS4 rather than saving three lives. Hence, in order to escape the second problem that actualism faces, we need to also posit a possibilist obligation that doesn't prescribe bad behavior.

Now, the *Gamer Gabi* case also helps illuminate the somewhat non-standard action-guiding notion we are appealing to in this paper, which concerns *minimizing wrongdoing in light of one's*

present circumstances.¹⁰ According to *Simple Hybridism**, Gabi has an actualist* obligation to purchase the PS4 today since, of all the act-sets under her present intentional control, performing that act-set would result in the best outcome. But Gabi also has a possibilist obligation to save the three lives tomorrow since that is the best outcome that she can bring about. Now, since fulfilling her actualist* obligation precludes fulfilling her possibilist obligation, we need to ask which obligation is action-guiding in the sense stipulated above. The actualist* obligation takes into account facts about an agent's rotten moral dispositions and prescribes the best option the agent can presently ensure that she performs. This allows the agent to minimize her wrongdoing in light of her present circumstances. So, the actualist* obligation is the action-guiding one. In other words, since saving the three lives tomorrow is not under Gabi's present intentional control, and since intending today to save the three lives tomorrow would result in *no* lives being saved, the thing for Gabi to do is what her actualist* obligation prescribes, viz. purchasing a PS4 today in order to prevent the three deaths from being extremely painful. With this understanding of our action-guiding notion in place, we now turn to *DOH*.

5. Dual Obligations Hybridism

DOH posits what we labeled in the previous section as an actualist* obligation. But, for reasons that will shortly become apparent, we will now refer to it as a synchronic obligation.¹¹ Now, in order to provide a more unifying account of obligations, instead of affirming a possibilist obligation, *DOH* posits the following obligation, which is defined *in terms* of a synchronic obligation:

At t an agent S has a ***diachronic obligation*** to ϕ at t' iff, if S were to have fulfilled all of S 's synchronic obligations at every time t^* such that $t^* < t$, then, given the moral character c S

¹⁰ Timmerman appeals to the same action-guiding notion in his [2015: 1525–7] when he defends the so-called 'Do Your Best' principle.

¹¹ As the reader will see shortly, *DOH* actually distinguishes between two kinds of synchronic obligations (and two kinds of diachronic obligations) in light of considerations about the proper parts of an obligatory act-set.

would have had as a result, if S were to in fact have c at t , then S would thereby have a synchronic obligation at t to ϕ at t' .¹²

By defining a diachronic obligation partly in terms of a synchronic obligation, *DOH* manages to posit a deeper connection between the two obligations in comparison to *Simple Hybridism* and *Simple Hybridism**. Moreover, the basic idea behind *DOH* is quite simple. An agent's synchronic obligation (which is action-guiding in the aforementioned sense) is a function of what she ought to do given her actual moral character, while an agent's diachronic obligation is a function of what she ought to do given how her moral character would have been (or might have been as indicated in footnote 12) if she had always done the right thing in the past. We want to be clear that when determining an agent's diachronic obligations, we hold fixed their skills, capacities, and the like. The only actual feature of the agent that is not being held fixed is the agent's moral character. Moreover, we explicitly remain neutral about which understanding of moral character is the correct one. For our purposes, we need only think of moral character in a very thin sense, viz. in terms of the extent to which an agent is disposed to sacrifice her own interests for the interests of other morally considerable beings.

One motivation for a diachronic obligation resembles the motivation for the similar notion of tracing. Many theories of moral responsibility account for an agent's derivative responsibility for some action by *tracing back* that action to some earlier action over which the agent had the requisite kind of control for grounding (non-derivative) responsibility. Roughly speaking, just as a proponent of tracing holds that an agent's past actions at least partly determine that agent's responsibility for some present

¹² There may be many ways to fulfill the same obligation, such that fulfilling that obligation in different ways makes a difference to the development of the agent's moral character. In that case, possibly, within the set of closest worlds in which an agent fulfills all of her synchronic obligations in the past up until the present moment, the worlds within that set vary with respect to the agent's moral character. Consequently, given the Stalnaker-Lewis semantics for counterfactuals, there may only be a fact about what an agent's moral character *might* have been like if that agent were to have fulfilled all of her synchronic obligations in the past. So, in such a case, which merely possible moral character is relevant to an agent's diachronic obligation? We suggest that the agent's actual moral character be replaced with the worst moral character in the set of closest worlds in which the agent fulfills all of her synchronic obligations in the past up until the present moment that is better than the agent's actual moral character. To be clear, there are various alternative proposals that might be offered. We wish to remain open to a variety of such proposals.

outcome, according to *DOH*, an agent's past actions likewise at least partly determine that agent's present diachronic obligations.

In order to see what *DOH* says with respect to *Evil Professor Smith*, let's suppose that, at t , (a)-ing and ($\sim a$)-ing are under Smith's intentional control, but X -ing is not. In that case, according to *DOH*, Smith has a synchronic (action-guiding) moral obligation to ($\sim a$). Moreover, let's stipulate that if Smith had fulfilled all of her synchronic obligations throughout her life up until this point, then given the moral character c Smith would have in such a case, if Smith were to in fact have c , then X -ing, Y -ing and Z -ing would be the act-sets under her present intentional control.¹³ In that case, Smith has a diachronic (non-action-guiding) obligation to X , which requires (a)-ing.

One final issue that must be considered is *DOH*'s stance on the proper subsets of (non-singleton) act-sets that an agent is obligated to perform.¹⁴ This issue can be posed in the form of a question: if an agent has an obligation to ϕ , and ϕ -ing requires ψ -ing, then what is the relationship between the agent and ψ -ing? In order to assess the deontic status of these proper subsets, we accept a version of the *Ought-Distributes-Over-Conjunction Principle*, which distinguishes between a dependent and a non-dependent moral obligation:

A moral obligation o is **dependent** iff an agent has o in virtue of some other moral obligation.

A moral obligation o is **non-dependent** iff o is not a dependent moral obligation.

With this distinction in hand, here is a plausible principle to which *DOH* is amenable:

¹³ We are stipulating that, as is the case for most persons, which act-sets are under Smith's present intentional control is partly determined by her moral character, and that Smith's moral character is partly determined by how she acts throughout her life; cf. Aristotle's *Nicomachean Ethics*, 1103b. For empirical evidence that our stipulation with respect to Smith is in fact true of most people, see, e.g., Baumeister and Tierney [2012: ch. 6].

¹⁴ For discussion of the problem concerning the sub-optimal consequences of performing proper subsets of act-sets that have the best consequences, see Bergström [1971], Castañeda [1973], Mendola [2006: ch. 2], and Gustafsson [2014].

The Dependence-Ought-Distributes-Over-Conjunction Principle At t an agent S has a dependent obligation to perform an act-set ψ iff (and because) at t S has a non-dependent moral obligation to perform an act-set ψ^* , such that ψ is a proper subset of ψ^* .

To illustrate how *DOH* subscribes to this principle, suppose as before that, at t , (a)-ing and ($\sim a$)-ing are under Smith's intentional control, but X -ing is not. Notice that since, at t , ($\sim a$)-ing is under Smith's intentional control, it follows from (2) that Y -ing is likewise under Smith's intentional control. In that case, at t Smith has a non-dependent synchronic (action-guiding) obligation to Y , and thus a dependent synchronic (action-guiding) obligation to ($\sim a$) (as well as ($\sim b$), (c), and (d)). Moreover, at t , Smith has a non-dependent diachronic (non-action-guiding) obligation to X , and thus has a dependent diachronic (non-action-guiding) obligation to (a) as well as (b).

To be clear, defending the above distinction between a dependent and non-dependent obligation is not the central concern of our paper. One may amend *DOH* to incorporate a different account of the deontic status of the proper subsets of (non-singleton) act-sets, which an agent is obligated to perform. For instance, one could adopt the strategy Portmore [2013] employs. Nevertheless, since *some* relationship must be posited between a (non-singleton) act-set an agent is obligated to perform and a proper part of that act-set, we only *suggest* understanding that relationship in the above manner.

We are finally in a position to provide our precise definition of *DOH*, which incorporates the dependent/non-dependent distinction:

Dual Obligations Hybridism (DOH) (I)–(IV) are true:

- (I) At t an agent S has a **non-dependent synchronic obligation** to ϕ at t' iff ϕ -ing at t' is an act-set under S 's intentional control at t , and what *would* happen if S were to ϕ at t' is better than what would happen if S were to perform any other incompatible act-set under S 's intentional control at t .
- (II) At t an agent S has a **non-dependent diachronic obligation** to ϕ at t' iff, if S were to have fulfilled all of S 's synchronic obligations at every time t^* such that $t^* < t$, then,

given the moral character c S would have had as a result, if S were to in fact have c at t , then S would thereby have a synchronic obligation at t to ϕ at t' .

- (III) At t an agent S has a (synchronic or diachronic) dependent obligation to perform an act-set ψ iff (and because) at t S has a (synchronic or diachronic) non-dependent moral obligation to perform an act-set ψ^* , such that ψ is a proper subset of ψ^* .
- (IV) At t , when an agent S 's synchronic and diachronic obligations conflict, S 's synchronic obligation(s) is (are) action-guiding.

Now that we see what *DOH* says with respect to *Evil Professor Smith*, we can show briefly how, like *Simple Hybridism*, *DOH* avoids the two *prima facie* problems that face actualism and the one that faces possibilism. We will stipulate once again that, at t , (a) -ing and $(\sim a)$ -ing are under Smith's intentional control, but X -ing is not.

First, unlike actualism, *DOH* does not allow an agent S to avoid incurring *any* moral obligation to ϕ , which S could easily fulfill, simply in virtue of S 's rotten moral character (unless one does not have such a character in virtue of failing to fulfill any of one's previous synchronic obligations; see section 7 below). According to *DOH*, although Smith does avoid incurring a (dependent and non-dependent) synchronic obligation to (a) , *DOH* disagrees with actualism that Smith lacks an obligation to (a) . Rather, Smith still has a dependent diachronic obligation to (a) in light of having a non-dependent diachronic obligation to X . So, while *DOH* permits S to avoid incurring a synchronic moral obligation to ϕ in virtue of having a rotten moral character, it's false on *DOH* that S can avoid incurring a moral obligation to ϕ *simpliciter* simply in virtue of having a rotten moral character.

Second, unlike actualism, *DOH* does not prescribe bad behavior, such that acting on such prescriptions renders S immune from moral criticism, even when S can easily refrain from such behavior. While *DOH* does hold that Smith is synchronically obligated to $(\sim a)$, Smith is still

diachronically obligated to (*a*). Thus, even if Smith performs ($\sim a$), Smith may still be criticizable for failing to fulfill her diachronic obligation under many plausible accounts of blameworthiness.¹⁵

Third, unlike possibilism, according to *DOH* it's false that *S* has an action-guiding obligation to ϕ even if *S* would perform a subsequent act-set that is deeply morally wrong, and worse than the act-set *S* would subsequently perform if *S* were to $\sim\phi$. For, *DOH* takes into account the fact that Smith would do something deeply wrong if she were to (*a*). So, according to *DOH*, Smith has an action-guiding, synchronic obligation to ($\sim a$) precisely because, *ex hypothesi*, the best act-set under Smith's intentional at *t*, is *Y*-ing, which includes ($\sim a$)-ing. Since *DOH* accounts for our seemingly conflicting intuitions in actualist/possibilist scenarios, and is immune from the three *prima facie* problems facing actualism and possibilism, there is good reason to adopt *DOH* over both actualism and possibilism.

One might worry that *DOH* gives rise to prohibition dilemmas in so far as our synchronic and diachronic obligations quite often conflict (since we don't in fact fulfill all of our synchronic obligations over our lifetime). We reply to this worry by noting that *DOH* is, at best, only committed to *self-imposed* moral dilemmas, which arise because of one's own wrongdoing in the past. *DOH* is not committed to *world-imposed* moral dilemmas, which can arise even when one has never performed a wrong action.¹⁶ The former kind of moral dilemma is less troublesome since it is more intimately linked with an agent's capacities throughout her life.¹⁷

6. Single Obligation Hybridism

Despite *DOH*'s advantages over actualism and possibilism, one might remain skeptical of *DOH* because it posits two obligations, because it might entail the possibility of self-imposed dilemmas, or

¹⁵ This is, of course, under the assumption that Smith is aware of all of the normatively relevant facts.

¹⁶ See Donagan [1984] and McConnell [1978] on self-imposed dilemmas.

¹⁷ *DOH* is not obviously committed to the possibility of self-imposed moral dilemmas. For, the following is necessarily true: *S* can fulfill her diachronic obligation, and if *S* were to do so, then *S* also would have fulfilled her synchronic obligation. This is because certain counterfactuals that are in fact true would be false if *S* were to fulfill *S*'s diachronic obligation.

because of the tracing-like factor of *DOH*'s diachronic obligation. We are sympathetic to these concerns. Since we are more confident in the success of *some* hybrid account that posits two moral 'oughts' than we are in *DOH* itself, we now wish to offer an alternative hybrid view that is immune from the above concerns: *Single Obligation Hybridism (SOH)*. According to *SOH*, there is only one moral obligation that is akin to the possibilist obligation, but which is explicitly not action-guiding in the aforementioned sense. Possibilists seemingly always take their obligation to be action-guiding, even if it's not explicitly stated. In fact, this very assumption is necessary for the standard objection against possibilism mentioned in section two to retain any force.¹⁸ So, to avoid confusion we will refer to the non-action-guiding obligation that *SOH* posits as a possibilist* obligation.

Next, according to *SOH*, there's also an action-guiding moral 'ought' that is not an obligation, which is akin to *DOH*'s synchronic obligation (or an actualist* obligation). This action-guiding moral 'ought' appeals to the notion of practical reason, whereby an agent *S* has most practical reason to ϕ iff *S* ought to ϕ in the action-guiding sense we have appealed to throughout this paper. In other words, if *S* has most practical reason to ϕ , then ϕ -ing *minimizes wrongdoing in light of one's present circumstances*. With this understanding in place, here is the second component to *SOH*:

Practically Best Principle (PBP) At *t* an agent *S* has most practical reason to ϕ at *t'* iff ϕ -ing at *t'* is under *S*'s intentional control at *t* and ϕ -ing at *t'* is either (i) identical to the non-dependent possibilist* obligation *S* has at *t*, (ii) a rationally permissible supererogatory act-set, or (iii) is the least rationally impermissible, all things considered, act-set under *S*'s intentional control at *t* (which may be either supererogatory, morally permissible or morally impermissible).¹⁹

¹⁸ Note, moreover, that no possibilist we know of has responded to the standard objection by denying that the possibilist obligation is action-guiding.

¹⁹ We are implicitly assuming here that moral rationalism is true. If moral rationalism is false, then *PBP* would require you to perform a rationally permissible, morally impermissible act-set instead of a rationally impermissible, morally permissible act-set whenever both options are under an agent's present intentional control.

There is an act-set that satisfies (iii) *iff* no act-set under S 's intentional control at t satisfies condition (i) or (ii).²⁰

An agent's possibilist* obligation and *PBP*'s prescription only come apart in certain crucial cases. This will be in any case where S has a possibilist* obligation to ϕ (which is an act-set that S can perform), yet ϕ -ing is not under S 's present intentional control, and so S ought, in the practical reason/action-guiding sense, to follow *PBP* by $\sim\phi$ -ing. This is why *PBP* explicitly states that what S has most practical reason to do is determined by condition (iii) when no act-set under S 's intentional control satisfies condition (i) or (ii). So, what is unique about *PBP* is that it picks out the act-set S has most reason to perform *among* the act-sets she can presently ensure that she will in fact perform. It's precisely because there are points in time where, in light of our moral shortcomings, we cannot ensure that we will fulfill our possibilist* obligation that we need an action-guiding principle like *PBP*.²¹ Lastly, just as we have applied the dependence/non-dependence distinction to obligations, the same holds, *mutatis mutandis*, for what an agent has most practical reason to perform.

We are now in a position to give a precise formulation of *SOH*:

Single Obligation Hybridism (SOH) (V)–(VII) are true:

- (V) At t an agent S has a ***non-dependent possibilist* obligation*** to ϕ at t' *iff* ϕ -ing at t' is part of the best maximally specific act-set that S can perform.
- (VI) At t an agent S has a dependent possibilist* obligation to perform an act-set ψ *iff* (and because) at t S has a non-dependent possibilist* moral obligation to perform an act-set ψ^* , such that ψ is a proper subset of ψ^* . Additionally, at t , S dependently has most practical reason to perform act-set ψ *iff* (and because) at t S non-dependently has most practical reason to perform an act-set ψ^* , such that ψ is a proper subset of ψ^* .

²⁰ In cases in which conditions (i)–(ii) are met, *ϕ -ing* will refer to the disjunction of the act-sets in conditions (i)–(ii).

²¹ One might wonder why *PBP* isn't identical to the right-hand side of the bi-conditional of our definition of a synchronic obligation. The primary reason is that agents sometimes have most practical reason to perform a rationally permissible act that is supererogatory or morally impermissible. Since *PBP* is not a moral obligation, it may be formulated in such a way that this important consequence of *SOH* is made explicit.

- (VII) At t an agent S has most practical reason to ϕ at t' iff ϕ -ing at t' is under S 's intentional control at t and ϕ -ing at t' is either (i) identical to the non-dependent possibilist* obligation S has at t , (ii) a rationally permissible supererogatory act, or (iii) is the least rationally impermissible, all things considered, act-set under S 's intentional control at t (which may be either supererogatory, morally permissible or morally impermissible). There is an act-set that satisfies (iii) iff no act-set under S 's intentional control at t satisfies condition (i) or (ii).

Unlike *DOH*, *SOH* is only committed to one type of obligation, and thus is not committed to the possibility of self-imposed dilemmas. Moreover, *SOH*'s possibilist* obligation does not employ a tracing-like factor, as is the case with *DOH*'s diachronic obligation, since *SOH* only posits one type of obligation.²² Furthermore, *SOH* retains the advantages that *DOH* has over both actualism and possibilism; *SOH* can account for the seemingly conflicting intuitions that arise in actualist/possibilist scenarios, and is also immune from the *prima facie* problems that arise for actualism and possibilism. To see how *SOH* works, let's consider *Evil Professor Smith* once more and assume once again that, at t , (a) -ing and $(\sim a)$ -ing are under Smith's intentional control, but X -ing is not. In that case, Smith has a non-dependent possibilist* obligation to X , and thus has a dependent obligation to (a) as well as (b) . Clearly, (a) -ing does not satisfy condition (i) or (ii). Moreover, Y -ing is the least rationally impermissible immoral act-set under Smith's intentional control at t . Thus, *PBP* entails that Smith non-dependently has most practical reason to Y , and thus dependently has most practical reason to $(\sim a)$ (as well as $(\sim b)$, (c) , and (d)).

Now, here is how *SOH* avoids the three aforementioned problems with actualism and possibilism respectively. First, unlike actualism, *SOH* does not permit an agent S to avoid incurring any moral obligation to ϕ , which S can easily fulfill, simply in virtue of S 's rotten moral character. Smith doesn't avoid incurring a non-dependent obligation to X , and thus doesn't avoid incurring a

²² Recall that by employing a tracing-like factor, *DOH* manages to posit a deep connection between its two types of obligations. However, since *SOH* is only committed to one type of obligation, *SOH* can dispense of a tracing factor without succumbing to the same *ad hoc* problem to which *Simple Hybridism* and *Simple Hybridism** are subject

dependent obligation to (a) . Second, *SOH* does not prescribe bad behavior, such that acting on such a prescription renders S immune from moral criticism, even when S can easily refrain from such behavior. After all, Smith's sole non-dependent obligation is her possibilist* obligation to X (which requires (a) -ing), and, on some plausible accounts of blameworthiness, Smith is blameworthy to the extent that she fails to fulfill her possibilist* obligation.²³ Finally, unlike possibilism, according to *SOH* S lacks an action-guiding obligation (or an action-guiding 'ought' in line with *PBP*) to ϕ even if S would perform a subsequent act-set that is deeply morally wrong, and worse than the act-set S would subsequently perform if S were to $\sim\phi$. For, *SOH* takes into account the fact that Smith would do something deeply wrong if she were to (a) . Thus, *SOH* entails that Smith non-dependently has most practical reason to Y , and thus dependently has most practical reason to $(\sim a)$.

7. Further Remarks on *DOH* and *SOH*

We have seen how *SOH* retains the advantages *DOH* has over actualism and possibilism, while also avoiding the aforementioned concerns with *DOH*. As a result, one may be inclined to prefer *SOH* to *DOH*. However, we will now investigate one substantive difference between *DOH* and *SOH* that concerns the obligations of a so-called instant agent, i.e., an agent that spontaneously comes into existence with mature cognitive faculties. Which hybrid view one prefers will depend, in part, on one's views about instant agents' obligations.

In comparison to *DOH*, *SOH* is more demanding of instant agents. To see this, consider *Evil Professor Smith*, but with the modification that Smith is an instant agent with no history. Call her Instant-Smith. Now, suppose that, at t , (a) -ing and $(\sim a)$ -ing are under Smith's intentional control, but X -ing is not. Additionally, assume that if Smith had fulfilled all of her synchronic obligations throughout her life up until t , then given the moral character c Smith would have had, if Smith were to in fact have c , then X -ing, Y -ing and Z -ing would be under her present intentional control.

²³ This is, of course, under the assumption that Smith is aware of all of the normatively relevant facts.

According to *DOH*, since Instant-Smith just came into existence, and thus did not have any obligations to fulfill prior to t , Instant-Smith has both a non-dependent synchronic obligation and a non-dependent diachronic obligation to Y . For, these two kinds of obligations cannot come apart with respect to an agent's very first obligation. By contrast, *DOH* entails that Smith has a non-dependent synchronic obligation to Y and a non-dependent diachronic obligation to X . This is because Smith had previous obligations to fulfill, such that if Smith had fulfilled all of her previous obligations, then, *ex hypothesi*, the moral character c that Smith would have had in that case is such that, at t , if Smith were to in fact have c , then X -ing would be under her intentional control. In contrast to *DOH*, *SOH* does not render different verdicts for Smith and Instant-Smith; both Smith and Instant-Smith have a non-dependent possibilist* obligation to X , but non-dependently have most practical reason to Y .

So, *DOH* affirms, while *SOH* denies, that Instant-Smith avoids incurring an obligation at t to X in virtue of the moral character that she has at "birth." One might resonate with *DOH*'s verdict since *DOH* takes into account the fact that, unlike Smith, Instant-Smith couldn't have done anything to develop a better moral character than the one she has. Alternatively, one might resonate with *SOH*'s verdict that both Smith and Instant-Smith have a non-dependent possibilist* obligation to X . For, even though Instant-Smith came into existence with a rotten moral character—a character she is presumably not blameworthy for having—it's still true that, at t , Instant-Smith can X in the same way that, at t , Smith can X . It's also stipulated that, at t , the type of control that Smith and Instant-Smith have over their actions is identical. Thus, there are a series of basic act-sets that both Smith and Instant-Smith can perform that are under their intentional control at various times that would result in them X -ing, viz. (a)-ing at t_1 and then (b)-ing at t_2 (see footnote 7). This might seem sufficient to generate a (non-dependent) moral obligation for both Smith and Instant-Smith to X .

8. Conclusion

Actualist/possibilist scenarios seem to generate conflicting intuitions. The only way to account for both intuitions is by positing a dual 'ought' view. *Hybridism*—the disjunction of *DOH* and *SOH*—not

only accounts for these intuitions, but also manages to be immune from the *prima facie* problems that face actualism and possibilism. While *Hybridism* fares better than any of its competitors, we leave it open to the reader as to whether *DOH* or *SOH* is the more plausible position. There are also other views in the vicinity of *DOH* and *SOH* that ought to be taken seriously. Still, we are optimistic that *Hybridism*, or something close enough, is the correct view concerning the nature of moral obligations.²⁴

References

- Baker, Derek 2012. Knowing Yourself—And Giving Up on Your Own Agency in the Process, *Australasian Journal of Philosophy* 90/4: 641–56.
- Baumeister, Roy F., and John Tierney 2012. *Willpower: Rediscovering the Greatest Human Strength*, New York: Penguin Books.
- Bergström, Lars 1971. Utilitarianism and Alternative Actions, *Noûs* 5/3: 237–52.
- Bykvist, Krister 2002. Alternative Actions and the Spirit of Consequentialism, *Philosophical Studies* 107/1: 45–68.
- Carlson, Erik 1999. Consequentialism, Alternatives, and Actualism, *Philosophical Studies* 96/3: 253–68.
- Castañeda, Héctor-Neri 1973. A Problem for Utilitarianism, *Analysis* 28/4: 141–42.
- Donagan, Alan 1984. Consistency in Rationalist Moral Systems, *The Journal of Philosophy* 81/6: 291–309.
- Feldman, Fred 1986. *Doing the Best We Can: An Essay in Informal Deontic Logic*, Dordrecht: Riedel.
- Goldman, Holly S. 1976. Dated Rightness and Moral Imperfection, *Philosophical Review* 85/4: 449–87.
- Goldman, Holly S. 1978. Doing the Best One Can, in *Value and Morals*, ed. Alvin Goldman and Jaegwon Kim, Dordrecht: D. Reidel Publishing Company. 185–214.
- Greenspan, Patricia S. 1978. Oughts and Determinism: A Response to Goldman, *Philosophical Review* 87/1: 77–83.
- Gustafsson, Johan E. 2014. Combinative Consequentialism and the Problem of Act Versions, *Philosophical Studies* 167/3: 585–96.
- Jackson, Frank, and Robert Pargetter 1986. Ought, Options, and Actualism, *Philosophical Review* 95/2: 233–55.
- McKinsey, Michael 1979. Levels of Obligation, *Philosophical Studies* 35/4: 385–95.
- McConnell, Terrance 1978. Moral Dilemmas and Consistency in Ethics, *Canadian Journal of Philosophy* 8/2: 269–87.
- Mendola, Joseph 2006. *Goodness and Justice*, New York: Cambridge University Press.
- Portmore, Douglas W. 2011. *Commonsense Consequentialism*, Oxford: Oxford University Press.
- Portmore, Douglas W. 2013. Perform Your Best Option. *The Journal of Philosophy* 110/8: 436–59.
- Portmore, Douglas W. Forthcoming. Consequentialism and Coordination: How Consequentialism

²⁴ We are grateful to (references deleted). We are also grateful to two anonymous referees for this journal for many extensive, helpful comments that in turn improved the paper significantly.

- Has an Attitude Problem, in *Consequentialism: New Directions, New Problems?*, ed. Christian Seidel. Oxford: Oxford University Press.
- Ross, Jacob 2013. Actualism, Possibilism, and Beyond, in *Oxford Studies in Normative Ethics vol. 2*, ed. Mark Timmons. Oxford: Oxford University Press: 74–97.
- Sobel, J. Howard 1976. Utilitarianism and Past and Future Mistakes, *Noûs* 10/2: 195–219.
- Timmerman, Travis 2015. Does Scrupulous Securitism Stand Up to Scrutiny? Two Problems for Moral Securitism and How we Might Fix Them, *Philosophical Studies* 172/6: 1509–28.
- Vorobej, Mark 2000. Prosaic Possibilism. *Philosophical Studies* 97/2: 131–36.
- Woodard, Christopher 2009. What’s Wrong with Possibilism. *Analysis* 69/2: 219–26.
- Wedgwood, Ralph 2009. Against Actualism, PEA Soup, posted on September 11, 2009. Accessed on May 16th, 2014 from <http://peasoup.typepad.com/peasoup/2009/09/againstaactualism.html>
- Zimmerman, Michael 1996. *The Concept of Moral Obligation*, New York: Cambridge University Press.
- Zimmerman, Michael 2006. The Relevant Risks to Wrongdoing, in *The Good, the Right, Life And Death: Essays in Honor of Fred Feldman*, ed. Kristopher McDaniel, Jason R. Raibley, Richard Feldman, and Michael J. Zimmerman, Burlington: Ashgate. 151–72.