



ACTUALISM HAS CONTROL ISSUES

BY YISHAI COHEN AND TRAVIS TIMMERMAN

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 10, No. 3 | OCTOBER 2016

URL: WWW.JESP.ORG

COPYRIGHT © YISHAI COHEN AND TRAVIS TIMMERMAN 2016

Actualism Has Control Issues

Yishai Cohen and Travis Timmerman¹

Whether an act is morally permissible or impermissible depends upon the alternative acts available to an agent. But what exactly are the *relevant* alternative acts available to an agent? The answer is far from obvious once we consider facts about how an agent *would* freely act in various situations. To get a grip on this abstract issue, consider the following concrete case:

The Gig: Brandi has been invited by her friend, Chad, to attend his musical gig at a local bar. Brandi can easily decide to attend the gig, and then decide at the gig to be supportive of Chad, which would be the best outcome. Unfortunately, Brandi detests Chad's music. Consequently, Brandi would not in fact decide to be supportive of Chad if she decided to attend his gig due to being irritated with Chad's performance – even though she *could* decide at the gig to be supportive. Since Chad would be deeply hurt, this would be the worst outcome. Brandi could alternatively decide not to attend Chad's gig, which would be better than the worst outcome, yet worse than the best outcome.

To be sure, Brandi *can* decide to attend the gig, and once there decide to be supportive of Chad. However, suppose that that is just not what Brandi *would* do if she were to attend. Here is the tricky philosophical question: is Brandi morally obligated to accept or decline the invitation to Chad's gig?

According to actualism, Brandi is obligated to decline the invitation because what would actually happen if Brandi declines is better than what would actually happen if she accepts (Goldman 1976; Sobel 1976; Jackson and Pargetter 1986). According to possibilism, Brandi is obligated to accept the invitation because accepting the invitation is part of the best act-set Brandi can perform over the course of her life (Greenspan 1978; Feldman 1986; Zimmerman (1996: § 6)). So the actualist affirms (and the possibilist denies) that facts about what agents would freely do in certain circumstances at least partly determine that agent's moral obligations.²

Various detailed accounts of both actualism and possibilism are on offer. Which particular definition one prefers makes no difference to our argument against actualism. We settle with the following formulations of each view:³

¹ This paper is the product of full and equal collaboration between its authors.

² This does not commit the actualist to the position that the truth of some counterfactual *c* concerning what an agent would freely do co-varies with some obligation *o* obtaining. Rather, the actualist is committed to the position that the mere truth of *c* can (but does not always) make a difference to *o*'s obtaining. By contrast, the possibilist holds that the mere truth of *c* cannot make a difference to whether *o* obtains.

³ Since the actualism/possibilism debate concerns obligations, the actions being considered are restricted to non-supererogatory ones.

Actualism: At t an agent S morally ought to ϕ at t' iff, at t , 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 morally ought to ϕ at t' iff, at t , ϕ -ing at t' is part of the best maximally specific act-set⁴ that S can perform.⁵

As is standard in the literature, throughout this paper we remain neutral with respect to the proper analysis of “can,” assuming that there is such an analysis. This is unproblematic since the actualism/possibilism debate cuts across issues concerning the proper analysis of “can.” Actualists and possibilists are concerned with the relationship between an agent’s free actions and their moral obligations. They are *not* concerned with what is required for an agent to have the ability to perform an action. Hence, as the reader will see, our argument against actualism will go through, irrespective of the account of “can” that the actualist might adopt.

Now, there are two prominent arguments against actualism. First, actualism permits agents to avoid incurring moral obligations, which they could (at the present moment) easily ensure that they will fulfill, simply in virtue of having rotten moral dispositions (Jackson and Pargetter (1986: 240); Timmerman (2006b: 156); Portmore (2011: 207); Baker (2012: 642-43); Timmerman (2015: 1512); Timmerman and Cohen (2016: 674)). Second, actualism prescribes bad behavior, even when the agent can (at the present moment) easily ensure that they refrain from such behavior (Wedgwood (2009); Ross (2012: 74-75); Timmerman (2015: 1513); Timmerman and Cohen (2016: 674); Cariani (forthcoming: 13-14)).

We propose a new argument against actualism. Actualism makes a morally irrelevant distinction between counterfactuals that do and do not at least partly determine an agent’s moral obligations, given that an agent sometimes has the same kind of control over the truth-value of both kinds of counterfactuals. We then offer a revised version of actualism that avoids this morally irrelevant distinction and that also has further advantages.

1. The Case

In order to set up our argument against actualism, it will be helpful to consider the following case:

⁴ At t , a maximally specific act-set that an agent can perform is one that extends from t to the last moment of time at which the agent can perform an act.

⁵ In any actual case, there will be equally good maximally specific act-sets that an agent can perform. Accordingly, possibilism says that an agent is morally obligated to fulfill one of the equally best maximally specific act-sets that she can perform. So, according to a more precise definition of possibilism, an act is obligatory iff that act is a member of every equally best maximally specific act-set that the agent can perform. An act is permissible iff that act is a part of at least one of the equally best maximally specific act-sets that the agent can perform.

Sadistic Suzy is a sadistic doctor who is responsible for the well-being of Sick Sid. At t_0 , Suzy can do the following:

- Completely cure Sid by giving him two shots of the complete cure (one at t_1 and one at t_2).
- Partially cure Sid by giving him two shots of the partial cure (one at t_1 and one at t_2).
- Let Sid die by giving him only one shot, giving him mixed shots or giving him no shots.

If Sid is completely cured, he will live a long and happy life for n years. If Sid is partially cured, he will live for $1/2n$ years. Under any other act-set that Suzy can perform, Sid dies immediately thereafter at t_3 . Now, take whatever account of control over an action that is required for an agent to have an obligation to perform that action, and suppose that, at t_0 , Suzy has *that* kind of control over each of the above act-sets. Next, let us suppose that the following counterfactuals are true:

- (i) If Suzy were to deliberate at t_0 about what to do about Sid, Suzy would give Sid no shot at t_1 .
- (ii) If Suzy were to give Sid no shot at t_1 , Suzy would give Sid no shot at t_2 (and Sid would immediately die).
- (iii) If Suzy were to give Sid a shot of the partial cure at t_1 , Suzy would give Sid a shot of the partial cure at t_2 (and Sid would be partially cured).
- (iv) If Suzy were to give Sid a shot of the complete cure at t_1 , Suzy would give Sid no shot at t_2 (and Sid would immediately die).

At t_0 , Suzy is in fact deliberating about what to do about Sid. This fact and the truth of (i) entail that Suzy gives Sid no shot at t_1 . Suzy's giving Sid no shot at t_1 and the truth of (ii) entail that Suzy gives Sid no shot at t_2 . Hence, Sid dies immediately thereafter at t_3 .

The actualist holds that the truth of (ii)-(iv) at least partly determines what obligations Suzy has at t_0 . For according to actualism, at t_0 , Suzy is obligated to give Sid a shot of the partial cure at t_1 because what would happen if Suzy were to do that (Sid would be partially cured) is better than what would happen if Suzy were not to give Sid a shot of the partial cure (Sid would immediately die). By contrast, the possibilist holds that, at t_0 , Suzy is obligated to give Sid a shot of the complete cure at t_1 since doing so is, *ex hypothesi*, part of the best act-set that Suzy can perform, viz., giving Sid both shots of the complete cure. Now, since Suzy in fact does not give Sid a shot at t_1 , both the actualist and the possibilist agree that, at t_1 , Suzy fails to fulfill a moral obligation, albeit different moral obligations for different reasons.

2. The Problem

Before introducing our proposed problem for actualism, we first need to gain a deeper understanding of the disagreement between actualists and possibilists. More specifically, we need to turn our attention to the relationship between the truth-value of certain counterfactuals and an agent's abilities. If an agent S has the ability to do otherwise, it follows that S has a kind of control

over the truth-value of certain counterfactuals, as we will now illustrate by looking at the *Sadistic Suzy* case.

Recall that the truth of (i) and the fact that, at t_0 , Suzy is deliberating about what to do about Sid jointly entail the following: Suzy does not give Sid a shot at t_1 . Nevertheless, at t_0 , Suzy can give Sid a shot at t_1 . Moreover, if Suzy were to give Sid a shot at t_1 , then (i) would be false and the following counterfactual that is in fact false would be true:

- (v) If Suzy were to deliberate at t_0 about what to do about Sid, she would give Sid a shot at t_1 .

So, since Suzy can give Sid a shot at t_1 , and “Suzy gives Sid a shot at t_1 ” entails (v), and (v) is false, it follows that Suzy has a kind of counterfactual power over the truth-value of both (i) and (v). Simply put, Suzy can do something, such that if she were to do it, then (i) would be false and (v) would be true.

Now, here is where actualists and possibilists disagree. Although the truth of (i) and the antecedent of (i) entail that Suzy does not give Sid a shot at t_1 , the possibilist maintains that, at t_0 , Suzy is obligated to do other than what she will in fact do, i.e., give Sid a shot of the complete cure at t_1 . So the truth-value of (i) does not even partly determine what, at t_0 , Suzy is obligated to do at t_1 . Likewise, possibilists also think that although the truth of (i), (ii) and the antecedent of (i) jointly entail that Suzy does not give Sid a shot at t_2 , Suzy is nevertheless obligated (at t_0) to do otherwise at t_2 . So the truth-value of (ii) similarly does not even partly determine what, at t_0 , Suzy is obligated to do at t_1 . The same goes for (iii) and (iv). The takeaway point is that possibilists hold that only facts about what Suzy *can* freely do, rather than facts about what Suzy *would* freely do, at least partly determine Suzy’s obligations.

The actualist agrees with the possibilist that, although the truth of (i) and the antecedent of (i) entail that Suzy does not give Sid a shot at t_1 , Suzy is still obligated (at t_0) to do other than what she in fact will do. More specifically, at t_0 Suzy is obligated to give Sid a shot of the partial cure at t_1 , even though this is not what she will in fact do. This is precisely because, at t_0 , Suzy can do otherwise at t_1 . However, unlike the possibilist, the actualist holds that, in order to determine what, at t_0 , Suzy ought to do at t_1 , it does not matter what Suzy *can* (at t_0) do at t_2 . Rather, what matters is what Suzy would do at t_2 were she to do something or other at t_1 . So the actualist maintains that it does not matter that, at t_0 , Suzy can do something that would render (ii) false, even though it does matter that, at t_0 , Suzy can do something that would render (i) false. We find this distinction that actualism makes to be morally irrelevant.

To emphasize, it is stipulated in *Sadistic Suzy* that the kind of control Suzy has, at t_0 , over all act-sets she can perform from t_1 to t_2 is the same. So why does the fact that Suzy would not give Sid a shot of the complete cure at t_2 allow her to avoid incurring an obligation at t_0 to give Sid a shot of the complete cure at t_1 ? After all, the fact that Suzy would not give Sid a shot of the

complete cure at t_1 does not similarly allow her to avoid incurring an obligation at t_0 to give Sid a shot of the complete cure at t_1 . Thus, this distinction that actualism makes between different kinds of counterfactuals appears to be morally irrelevant. Indeed, generalizing from *Sadistic Suzy*, we can see that actualism entails an implausible thesis. In colloquial terms, actualism entails that some (future) acts under an agent's control do partly determine an agent's present obligations, while other (future) acts under an agent's control in the exact same sense do not even partly determine an agent's present obligations. In order to state this implausible thesis more precisely and formally, we must first give names to the following two types of counterfactuals:

An **early** counterfactual concerns what S would do at t_1 given the circumstances that S is in at t_0 .

A **late** counterfactual concerns what S would do at t_2 given the circumstances that S is in at t_1 .

With these two types of counterfactuals in place, we have shown that actualism entails the following implausible thesis:

Actualism's Consequence: For any obligation o agent S has at t_0 to ϕ at t_1 , and for any early counterfactual, and for any late counterfactual, o is not even partly determined by the early counterfactual(s), but is at least partly determined by the late counterfactual(s).

Actualism's Consequence demarcates between counterfactuals that do and do not determine S 's obligations, even in cases in which, at t_0 , S has the same kind of control over what S does both at t_1 and t_2 (and thus has the same kind of control over the truth-value of both kinds of counterfactuals mentioned in *Actualism's Consequence*). In other words, *Actualism's Consequence* is committed to what we think is a morally irrelevant distinction between two different kinds of counterfactuals. At least, it is a morally irrelevant distinction with respect to what determines an agent's obligations. So, since actualism entails *Actualism's Consequence*, and *Actualism's Consequence* is implausible, actualism is itself implausible.

We wish to emphasize that our critique of actualism is neutral with respect to the proper analysis of "can." We take this to be a virtue, not a bug, of our criticism precisely because whatever account of "can" the actualist wishes to adopt, she will nevertheless be committed to *Actualism's Consequence*, which remains implausible under any viable account of "can." In the next section, we offer a way for the actualist to avoid *Actualism's Consequence* that likewise remains neutral with respect to the proper analysis of "can."

3. The Solution

The way out for the actualist is to accept a revised version of actualism (or a view that retains the spirit of actualism) that does not entail *Actualism's Consequence*. We suggest that the way to do this is by demarcating counterfactuals on the basis of a kind of control an agent can have over their truth-value. We will illustrate how such a version of actualism can avoid *Actualism's Consequence* in two steps. First, we employ the following notion of intentional control:

ϕ -ing at t' is under an agent S 's *intentional control* at t iff, at t , S can⁶ form a certain set of intentions, such that if S were to form such intentions, then S would ϕ at t' .⁷

We will show that appealing to a notion of control, such as the one above, will render one's view immune from the aforementioned problem. But before we show this, there is a second step in our revised version of actualism. Instead of putting forward a view that only considers whether an agent will ϕ or $\sim\phi$, we need a view that explicitly considers all of the actions that are incompatible with ϕ -ing. In order to see why, consider the following case that is adapted from Wedgwood (2009) and Ross (2012: 75):

Deadly Dexter enjoys killing people for fun, but typically resorts to only killing criminals. Dexter is presently in a room with three innocent individuals: Maria, Vince and Joey. Dexter can kill any set of these individuals, or refrain from killing anyone. Moreover, Dexter has the morally relevant kind of control over all of these actions. However, since he is presently in a killing mood, Dexter would kill both Joey and Maria if he were to refrain from killing Vince.

Actualism renders the absurd verdict that Dexter ought to only kill Vince, even though he can refrain from killing all three of them. So it is for this reason that we should not say that an agent S is morally obligated to ϕ just in case what would happen if S were to ϕ is better than what would happen if S were to $\sim\phi$. Instead, one should say that, of all of the actions S can perform at a time t , S ought to perform the action at t that would result in the best

⁶ To repeat, we remain neutral with respect to the proper analysis of "can." Consequently, we neither affirm nor deny the position that, at t , S can ϕ at t' iff ϕ -ing at t' is under S 's intentional control at t . Irrespective of what the actualist maintains in regard to this matter, actualism will still be committed to *Actualism's Consequence*.

⁷ As is standard in the literature, we are first and foremost concerned with objective obligations rather than subjective ones (c.f. Zimmerman 2006a). Roughly, *objective obligations* are those determined by all of the normatively relevant facts, including the normatively relevant facts of which the agent is unaware. *ICA* entails that agents can be objectively obligated to perform acts they do not know are under their intentional control. Blameworthiness, however, tracks *subjective obligations*, those that an agent ought to perform relative to the evidence in her epistemic ken. Readers who deny the existence of objective obligations can amend *ICA*, such that it concerns subjective obligations.

outcome (c.f. Cariani (forthcoming: 13-14)). The same holds, *mutatis mutandis*, for a view that employs the notion of intentional control.

Now that these two revisions to actualism have been laid out, we are in a position to present the view that is superior to actualism, and, more specifically, is able to avoid *Actualism's Consequence*:

Intentional Control Actualism (ICA): At t an agent S morally ought⁸ 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 .⁹

ICA does not entail *Actualism's Consequence*. In fact, *ICA* is inconsistent with *Actualism's Consequence*. To illustrate, suppose we further stipulate in *Sadistic Suzy* that giving Sid either shot or giving Sid no shot at t_1 is under Suzy's intentional control at t_0 . However, it is not under Suzy's intentional control at t_0 to give Sid both shots of the complete cure. In that case, both *ICA* and actualism entail that, at t_0 , Suzy ought to give Sid a shot of the partial cure at t_1 . So, in this case, the truth of (iv) does at least partly determine Suzy's obligation at t_0 since the consequent of (iv) concerns an act-set that is not under Suzy's intentional control at t_0 . However, if we instead stipulate in *Sadistic Suzy* that any combination of giving or refraining from giving Sid shots is under Suzy's intentional control at t_0 , then *ICA* disagrees with actualism that, at t_0 , Suzy ought to give Sid a shot of the partial cure at t_1 . Instead, *ICA* holds that, at t_0 , Suzy is obligated to give Sid both shots of the complete cure (one at t_1 , another at t_2) since what would happen if Suzy were to perform this act-set is better than what would happen were she to perform any other act-set under her intentional control at t_0 . Consequently, according to *ICA*, counterfactuals (ii)-(iv) do not even partly determine any of Suzy's obligations at t_0 .

The lesson to be drawn from these cases is that, according to *ICA*, a necessary condition for a counterfactual to even partly determine an agent's

⁸ As the reader will see in section 4, we distinguish between two kinds of moral obligations (dependent and nondependent) in light of considerations about the proper parts of an obligatory act-set. Moreover, *ICA* only concerns nondependent obligations.

⁹ *ICA* is similar in certain respects to the view proposed by Holly Smith (1978: 202) as well as the revised views of Carlson (1995: 121-23) and Bykvist (2002: 61-64) in light of their criticisms of Smith's view. Both Carlson and Bykvist reject Smith's view since it employs a conditional analysis of abilities, and they both reject such an analysis. Since, as noted in n. 6, *ICA* is explicitly silent about the proper account of abilities, *ICA* is immune from this criticism.

The views of Carlson and Bykvist share (at least) one particularly noticeable difference with *ICA*. Carlson (1995: 81-82) and Bykvist (2002: 50-51) formulate their respective theories with the assumption that ϕ -ing is presently performable by an agent only if the agent can presently form an intention to ϕ . By contrast, *ICA* grants that an agent can have an objective obligation to ϕ that the agent can perform only by forming an intention to do something other than ϕ (assuming that the agent can form this intention). We welcome this consequence since the actualism/possibilism debate concerns objective obligations rather than subjective obligations. Only the latter type of obligation tracks moral responsibility.

present obligation is that the consequent of that counterfactual refer to an act-set (including singleton act-sets) that is not under an agent's present intentional control. This is precisely why *ICA* is inconsistent with *Actualism's Consequence*. Instead, *ICA* entails the following thesis:

ICA's Consequence: For any obligation o agent S has at t_0 to ϕ at t_1 , and for any early counterfactual, and for any late counterfactual, o is not even partly determined by the early counterfactual(s), but is at least partly determined by the late counterfactual(s) *only if the consequent of the late counterfactual refers to an act-set that is not under S 's intentional control at t_0 .*

ICA's Consequence is not implausible precisely because it demarcates those counterfactuals that at least partly determine an agent's obligations and those that do not on the basis of the kind of control the agent has. By contrast, *Actualism's Consequence* does no such thing. According to *Actualism's Consequence*, there need not be any difference with respect to the kind of control an agent has over certain act-sets in order for one counterfactual to at least partly determine an agent's obligation, and for another counterfactual to not even partly determine an agent's obligation.¹⁰

To be clear, we are not arguing for the truth of *ICA*. Nor are we arguing for the truth of any revised version of actualism that avoids *Actualism's Consequence*. Rather, we are only arguing that the best versions of actualism avoid *Actualism's Consequence*. Additionally, in order to avoid *Actualism's Consequence*, the actualist should accept a version of actualism that incorporates a kind of control that demarcates counterfactuals on this basis.

Besides *ICA*, there are alternative views that escape *Actualism's Consequence* in virtue of incorporating a kind of control that demarcates the relevant kinds of counterfactuals from one another. For instance, the securitist views of Doug Portmore and Jacob Ross, respectively, offer notions of con-

¹⁰ As shown in the previous section, possibilism is immune from this problem by treating counterfactuals (ii)-(iv) in a more consistent manner, viz., by holding that none of the counterfactuals (ii)-(iv) even partly determines what obligations Suzy has at t_0 . However, we think that possibilism is problematic insofar as it also does not demarcate between counterfactuals that determine an agent's moral obligations and those that do not on the basis of the kind of control an agent has. This gives rise to the main objection to possibilism in the literature, viz., that possibilism says that in certain circumstances an agent S ought to ϕ even if S would perform a subsequent act that is deeply morally wrong, and worse than what S would subsequently perform if S were to $\sim\phi$. For discussion of this objection in the literature, see Goldman (1976: 469-70), Almeida (1992: 461-62), Woodard (2009: 219-20) and Portmore (2011: 211).

This consequence is counterintuitive at least in cases in which the act-set that S is obligated to perform is not under S 's intentional control (or some other kind of control sufficiently similar to intentional control). For instance, suppose once more that, in *Sadistic Suzy*, giving Sid either shot or giving Sid no shot at t_1 is under Suzy's intentional control at t_0 . But it is not under Suzy's intentional control at t_0 to give Sid both shots of the complete cure. Accordingly, possibilists hold that, at t_0 , Suzy is obligated to give Sid a shot of the complete cure at t_1 , despite the fact that, no matter what Suzy does (or intends to do) at t_0 , she will not give Sid a shot of the complete cure.

trol that the actualist could adopt in order to avoid *Actualism's Consequence*. Nevertheless, we think that intentional control is the morally relevant sense of control in this debate.¹¹ Hence, before moving on to the next section, we will defend this claim.

Here is the notion of control to which Portmore appeals:

[S]omething is *scrupulously* securable by an agent if and only if there is both some set of intentions and some set of *permissible* background attitudes such that if she were to have both those intentions and those permissible background attitudes, that something would be brought about (Portmore 2011: 165).

While Portmore's (2011: 167-69) official version of scrupulous securability is more complicated, the above account suffices for our purposes. So, roughly speaking, Portmore holds that an agent *S* ought to ϕ iff performing the best act-set that is scrupulously securable by *S* requires ϕ -ing. Accordingly, it is possible for an agent *S* to be morally obligated to ϕ when ϕ -ing requires having a different set of background attitudes than the ones that *S* in fact has. These attitudes include such mental states as desires and beliefs (Portmore 2011: 167).

While we lack the space to fully describe Ross' position, the important thing to note for our purposes is that, according to Ross' notion of a directly securable option, it is likewise possible for an agent *S* to be morally obligated to ϕ when ϕ -ing requires having a different set of background attitudes than the ones that *S* in fact has (see Ross (2012: 82, 91, 94-95)).

With this understanding of these two notions of control, we offer two objections to each notion of control being the morally relevant one. First, an agent *S* cannot be morally obligated to ϕ if ϕ -ing requires *S* to presently have different desires or beliefs from the ones that *S* presently has. For our desires and beliefs are arguably not directly up to us in the same way that basic mental actions (such as decisions) are up to us. To be clear, we *can* change our desires and beliefs over time by performing various actions, such as reevaluating our evidence, reconsidering the intrinsic value of our desires and so on (Alston 1989). But we arguably cannot directly alter our beliefs or desires without first performing an action of some kind. Hence, the morally relevant type of control that is to be incorporated into the view that solves the actualism/possibilism debate is one that does not require an agent to have a different set of background attitudes than the ones the agent in fact has. The notion of intentional control satisfies this *desideratum*.¹²

Second, as a consequence of Portmore's and Ross' views not holding fixed the agent's actual background attitudes, their versions of securitism

¹¹ We employ the notion of intentional control in Timmerman and Cohen (2016).

¹² This *desideratum* is consistent with the so-called tracing condition for responsibility, according to which one can be derivatively blameworthy for one's present desires or beliefs in virtue of previously performing an action (or set of actions) for which one is non-derivatively blameworthy.

succumb to a problem to which possibilism succumbs: it generates action-guiding obligations to perform acts that, if performed, will result in the worst possible outcome. To see the problem, it will be helpful to consider a variation of one of Ross' (2012: 83) own cases.

General Ripper 2: General Jack D. Ripper is in the following situation. It is now noon. At 12:01, he will face the choice of pressing or not pressing a button. At 12:02, he will face that same choice. If he presses the button at both times, then nothing happens. If he refrains from pressing the button at both times, then Innocent Idris will receive a momentary electric shock at 12:03. If he presses the button at one of the times, but not the other, then Innocent Idris will receive a momentary painful electric shock at 12:03 and the world will be destroyed at 12:04.

Ripper happens to dislike Innocent Idris (for no good reason) and wants him to be shocked to such an extent that he prefers any outcome in which Idris is shocked over any outcome in which he is not. As Ross and Portmore would grant, this is an impermissible background attitude. Now, suppose that, at noon, no matter what Ripper intends to do at 12:02, he will not push the button at 12:02. After all, Ripper wants to ensure that Idris is shocked, even if this results in the destruction of the world, though he prefers to shock Idris without destroying the world. Now, the following counterfactuals are true.

- (vi) If Ripper were not to push the button at 12:01, he would not push the button at 12:02 (and Idris would receive an electric shock, but the world would remain intact).
- (vii) If Ripper were to push the button at 12:01, he would not push the button at 12:02 (and Idris would receive an electric shock and the world would be destroyed).

Let us also suppose that if Ripper had permissible background attitudes at noon, then his pushing the button at 12:01 would entail that he also pushes the button at 12:02, thereby saving the world and preventing Innocent Idris from being shocked. Now here is the problem. According to Portmore's and Ross' securitist views, Ripper ought to push the button at 12:01, despite the fact that no matter what Ripper intends at noon to do, he will not push the button at 12:02. So securitism entails the implausible thesis that agents can have an action-guiding obligation to act in ways that, no matter what intentions they presently form, will result in the worst possible outcome. Elsewhere we have argued that this a problem to which possibilism succumbs (Timmerman and Cohen 2016). *ICA*, however, avoids this implausible consequence. According to *ICA*, Ripper ought to refrain from pushing the button at 12:01. To see why, note that, at noon, the only options under Ripper's intentional control are the following:

- (a) Refrain from pushing the button at 12:01 and refrain from pushing the button at 12:02.
- (b) Push the button at 12:01 and refrain from pushing the button at 12:02.

Since (a) would result in a much better outcome than (b), Ripper is required to (a), which requires refraining from pushing the button at 12:01. *ICA* thus seems to render the intuitively correct verdict. At the very least, we think *ICA* better captures actualist intuitions than securitism.

Before moving on to the next section, we wish to emphasize that, although we think that intentional control is the morally relevant sense of control for the purposes of this debate, we are not arguing for the truth of *ICA* or any other revised version of actualism. Rather, we are arguing that *Actualism's Consequence* is implausible, and that, in order to avoid *Actualism's Consequence*, the actualist must adopt a position like *ICA* that demarcates counterfactuals on the basis of a type of control an agent possesses. Moreover, the actualist could adopt a version of *ICA* that instead employs Portmore's or Ross' type of control in order to avoid *Actualism's Consequence*. However, for reasons just given, we think that the actualist should prefer our version of *ICA*. In the next section, we will note one further advantage that *ICA* enjoys over actualism. Then, in section 5, we will consider how the actualist might reply to our main argument.

4. A Further Advantage

Unlike *ICA*, actualism is not action-guiding in the following sense: for any act-set S can perform, actualism says whether S ought to perform that act-set. But many of these act-sets S ought to perform are incompatible with one another, and actualism does not say which obligation takes priority (Zimmerman (1996: 191-93); Carlson (1999: 261); Vessel (2009: 187-91)).¹³ To see this, recall the definition of actualism offered at the outset:

Actualism: At t an agent S morally ought at t to ϕ at t' iff, at t , 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' .

Notice that, on this definition of actualism, like any other, there are no restrictions on which act-sets can stand in for " ϕ " or what times can stand in for " t' ." So far we have seen that, in *Sadistic Suzy* where " ϕ " refers to giving Sid a shot of the partial cure at t_1 , actualism holds that, at t_0 , Suzy ought to ϕ because what would happen if Suzy were to give Sid a shot of the partial cure

¹³ Jennie Louise (2009: 330) considers this objection, but argues that it rests on a misunderstanding of actualism since, on actualism, it is always possible for an agent to act in such a way that she fulfills all of her obligations. Given Louise's implicit assumption that "can" is not coextensive with "intentional control," the following is true: while an agent *can* always act in such a way so as to fulfill all of her obligations, there are possible circumstances in which fulfilling all of her obligations is not under her *present intentional control*. So, in such circumstances, it still appears to be an open question as to what the agent should do in the action-guiding sense.

at t_1 (Sid would be partially cured) is better than what would happen if she were not to do this (Sid would immediately die).

But now suppose instead that “ ϕ ” refers to giving Sid both shots of the complete cure, and that “ t ” corresponds to t_1 - t_2 .¹⁴ When considering this act-set, actualism holds that, at t_0 , Suzy ought to give Sid both shots of the complete cure from t_1 to t_2 because what would happen if Suzy were to perform this act-set (Sid would live for n years) is better than what would happen if Suzy were not to perform this act-set (Sid would either live for $1/2n$ years or immediately die). In other words, although actualism holds that, at t_0 , Suzy is obligated to give Sid a shot of the partial cure at t_1 , Suzy is *also* obligated (at t_0) to do something that is incompatible with that, viz., to both give Sid a shot of the complete cure at t_1 and give Sid a shot of the complete cure at t_2 . So actualism is not action-guiding insofar as actualism (quite often) generates incompatible moral obligations, and furthermore is silent about which obligation takes priority.

By contrast, *ICA* is action-guiding; *ICA* does not entail that one can have conflicting obligations since there are restrictions on which act-sets can stand in for “ ϕ ” and which times correspond to “ t .” The only act-set that can stand in for “ ϕ ” is the one under an agent’s intentional control that, if performed, would bring about a better outcome than any of the other act-sets under the agent’s intentional control (“ t ” will thus correspond to the time required for performing that act-set). So, unlike actualism, *ICA* does not fail to be action-guiding insofar as *ICA* does not generate incompatible moral obligations.¹⁵

Perhaps this seems too quick. Since *ICA* does not restrict an agent’s obligations to the maximal sequence of acts under an agent’s intentional control, *ICA* might seem to generate incompatible obligations of the following sort: suppose once more that, at t_0 , any combination of giving or refraining from giving Sid shots is under Suzy’s intentional control. According to *ICA*, Suzy has an obligation to perform the act-set of giving Sid both shots of the complete cure (one at t_1 and one at t_2). This is because, of all the act-sets under Suzy’s intentional control at t_0 , performing this act-set would result in the best outcome. Now, one might wonder whether, according to *ICA*, Suzy is (at t_0) obligated to perform the singleton act-set of giving Sid a shot of the

¹⁴ Given that “ ϕ ” can refer to act-sets whose completion extends over a duration of time, “ t ” may similarly refer to a duration of time (such as the duration from t_1 - t_2), rather than only refer to a single moment of time.

¹⁵ By solving actualism’s action-guiding problem, *ICA* fails to render the verdict that Suzy has an obligation to give Sid a shot of the complete cure at t_1 and at t_2 (supposing it is not under Suzy’s intentional control to give Sid a shot of the complete cure at t_1 and at t_2). Depending on what account of “can” one thinks is relevant to moral obligations, this may seem to be a cost to *ICA* (cf. Vessel 2016). But this is, at best, a minor cost in comparison to all of the costs that plague actualism. Ultimately, we reject actualism in favor of a hybrid account in Timmerman and Cohen (2016). However, our goal in this paper is to explore how actualists can best develop their view in response to the central argument being raised in this paper.

complete cure at t_1 . It is natural to suppose that she is not because Sid would die if Suzy were to give Sid a shot of the complete cure at t_1 (barring, of course, an intention to give Sid both shots of the complete cure). Consequently, it seems impermissible for Suzy to intend to perform this singleton act-set.

So why is this not a case in which *ICA* prescribes incompatible obligations? The answer relies on the distinction between *dependent* and *nondependent* moral obligations, which we define in the following way.

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

A moral obligation m is **nondependent** iff it is not a dependent moral obligation.

The nondependent moral obligation is whichever act-set “ ϕ ” picks out, according to *ICA*. In the case being considered, Suzy’s nondependent moral obligation (at t_0) is to give Sid a shot of the complete cure at t_1 and at t_2 . Now, fulfilling this nondependent moral obligation requires that Suzy give Sid a shot of the complete cure at t_1 . Accordingly, Suzy has (at t_0) a dependent moral obligation to give Sid a shot of the complete cure at t_1 . It also requires her to give Sid a shot of the complete cure at t_2 . So Suzy likewise has (at t_0) a dependent moral obligation to give Sid a shot of the complete cure at t_2 . However, Suzy does not have a nondependent obligation to give Sid a shot of the complete cure at t_1 since intending to perform that singleton act-set would not result in the best outcome available to Suzy.¹⁶

Drawing this distinction keeps *ICA* consistent with the “ought distributes over conjunction” principle and prevents *ICA* from prescribing incompatible obligations.¹⁷ It also makes *ICA* uniquely action-guiding: an agent always ought to form an intention that will result in her fulfilling her nondependent moral obligation because, in doing so, she will always succeed in fulfilling all of her dependent and nondependent moral obligations. We will now consider and rebut an objection to our main argument.

5. An Objection

One might suspect that our challenge to actualism is merely a pseudo-problem, suggesting that the core commitment of actualism itself reveals the relevant difference between early and late counterfactuals. Here is a more detailed description of the sort of objection we have in mind. Recall the following two counterfactuals in *Sadistic Suzy*:

¹⁶ For simplicity’s sake, we are assuming in *Sadistic Suzy* that, for any act-set x under Suzy’s intentional control, Suzy can only succeed in x -ing by intending to x rather than by intending to perform some action distinct from x .

¹⁷ For further discussion of this vexing and underexplored issue, see e.g., Bergstrom (1971), Castañeda (1973), Mendola (2006) and Portmore (2013).

- (i) If Suzy were to deliberate at t_0 about what to do about Sid, Suzy would give Sid no shot at t_1 .
- (iv) If Suzy were to give Sid a shot of the complete cure at t_1 , Suzy would give Sid no shot at t_2 (and Sid would immediately die).

Actualism's Consequence entails that the truth of (i) does not even partly determine Suzy's t_1 -obligations (at t_0), but the truth of (iv) does at least partly determine Suzy's t_1 -obligations (at t_0). The actualist might argue that they have a perfectly good explanation for this: Suzy avoids incurring an obligation (at t_0) to give Sid a shot of the complete cure at t_1 because (iv) entails that what would happen if Suzy were to give a shot of the complete cure at t_1 is worse than what would happen if Suzy were not to do this. The truth of (i), on the other hand, does *not* entail that what would happen if Suzy were to give a shot of the complete cure at t_1 is worse than what would happen if Suzy were not to do this. In other words, (iv) is relevant to Suzy's t_1 -obligations (at t_0) since it entails facts about the relative values of the outcomes of her available t_1 -actions. By contrast, (i) is irrelevant to Suzy's t_1 -obligations (at t_0) since it does not entail such facts. Actualists might then attempt to rebut our argument by highlighting this difference and then digging in their heels.

While we think this objection fails, let us first be clear about the issues on which we agree. We agree that the consequences of an action can make a difference to what an agent is obligated to do. This position can be accepted by everyone in the actualist/possibilist debate. To illustrate, suppose that Ryuchi can either press button A, which will save 10 lives, or press button B, which will save only one of the 10 lives. Suppose, moreover, that nothing Ryuchi could do subsequent to pressing one of the buttons affects the number of lives that are saved. In that case, everyone in the actualist/possibilist debate can agree that Ryuchi ought to press button A precisely because of the consequences of pressing this button. What we disagree with is that the consequence c of some action a at least partly determines an agent's present obligations when c is identical to an action a^* by the agent, and that agent has the exact same type of control over a that she has over a^* . Now that this potential ambiguity is cleared up, we will explain why the aforementioned objection to our argument fails, in two parts.

First, while standard versions of actualism do provide a principled or non-ad-hoc method for determining which counterfactuals at least partly determine an agent's obligations, it does not follow that this method is morally relevant. To illustrate with an analogy, consider a speciesist who holds that only members of her species have a certain moral status, solely in virtue of being *Homo sapiens*. While such a speciesist employs a principled or non-ad hoc distinction between beings that do and do not have a certain moral status, this person is employing a morally irrelevant distinction. So the mere fact that actualism employs a non-ad-hoc method for demarcating between certain counterfactuals in no way entails that such a method marks a morally relevant distinction.

Second, we can see that this method is morally irrelevant by revisiting a problem for actualism that was highlighted in the *Deadly Dexter* case. Recall that actualism renders the intuitively incorrect verdict that Dexter ought to kill Vince. Actualism's verdict in this case is certainly governed by a non-ad-hoc principle: Dexter ought to kill Vince because, if he were to refrain from doing so, then he would kill both Joey and Maria. However, this principle is morally irrelevant precisely because Dexter possesses the same kind of control over refraining from killing all three individuals as he does over killing any set of them. Similarly, actualism's demarcation between counterfactuals (i) and (iv) is morally irrelevant given that actualism's distinction cuts across the kind of control an agent has over the actions she can perform. It is certainly true that (iv) entails facts about the relative values of the outcomes of her available t_1 -actions (at t_0), and that (i) does not entail such facts. But this does not get at the heart of our challenge to actualism, which concerns the kind of control an agent has over the actions she can perform.

We want to know why, at t_0 , Suzy's rotten moral disposition described in counterfactual (iv) allows her to avoid incurring moral obligations that her rotten moral disposition described in counterfactual (i) does not, given that, at t_0 , Suzy has (*ex hypothesi*) the exact same control over the truth-value of each counterfactual. After all, the actualist agrees that Suzy can be obligated to do other than what she will in fact do at t_1 given the truth of (i) (and her present circumstances). So why is Suzy's obligation to do something at t_1 dependent upon what Suzy would do thereafter since Suzy can still do other than what she would do thereafter?

Once it is stipulated that the type of control Suzy has, at t_0 , over the truth-value of counterfactuals (i) and (iv) is identical, there is compelling reason to hold that either both (i) and (iv) at least partly determine Suzy's t_1 -obligations (at t_0) or neither do. Unless the actualist provides a compelling reason to think that this line of reasoning is mistaken, the aforementioned objection fails. We thus conclude that the actualist should abandon her position in favor of *ICA* or a similarly structured view that is not committed to *Actualism's Consequence* and that retains *ICA's* aforementioned advantages.¹⁸

Yishai Cohen
Syracuse University
Department of Philosophy
yishai.a.cohen@gmail.com

Travis Timmerman

¹⁸ For helpful feedback on earlier drafts of this paper, we are grateful to Ben Bradley, Sean Clancy, Jamie Dreier, Nicole Dular, Kirsten Egerstrom, Karl Ekendahl, Andrew Forcehimes, Peter A. Graham, Matthew Koehler, Amy Massoud, Teresa Bruno Nino, Howard Nye, Hille Paakkunainen, Doug Portmore, Nate Sharadin, Luke Semrau, David Sobel, Aaron Wolf, Michael Zimmerman and our audience at the 2014 Pacific APA. We are also greatly indebted to the anonymous referees who read this paper and whose detailed and insightful comments significantly improved the final product.

Seton Hall University
Department of Philosophy
travis.timmerman@shu.edu

References

- Almeida, M. (1992) "The Paradoxes of Feldman's Neo-Utilitarianism," *Australasian Journal of Philosophy* 70(4): 455-68.
- Alston, W. (1989) "The Deontological Conception of Epistemic Justification," in *Essays in the Theory of Knowledge*, Ithaca, NY: Cornell University Press, pp. 115-52.
- Baker, D. (2012) "Knowing Yourself – and Giving Up on Your Own Agency in the Process," *Australasian Journal of Philosophy* 90(4): 641-56.
- Bergstrom, L. (1971) "Utilitarianism and Alternative Actions," *Noûs* 5(3): 237-52.
- Bykvist, K. (2002) "Alternative Actions and the Spirit of Consequentialism," *Philosophical Studies* 107(1): 45-68.
- Cariani, F. (forthcoming) "Consequence and Contrast in Deontic Semantics," *The Journal of Philosophy*.
- Carlson, E. (1999) "Consequentialism, Alternatives, and Actualism," *Philosophical Studies* 96(3): 253-68.
- _____. (1995) *Consequentialism Reconsidered*, Dordrecht: Riedel.
- Castañeda, H. (1973) "A Problem for Utilitarianism," *Analysis* 33: 141-42.
- Feldman, F. (1986) *Doing the Best We Can: An Essay in Informal Deontic Logic*, Dordrecht: Riedel.
- Goldman, H. (1976) "Dated Rightness and Moral Imperfection," *Philosophical Review* 85(4): 449-87.
- Greenspan, P. S. (1978) "Oughts and Determinism: A Response to Goldman," *Philosophical Review* 87(1): 77-83.
- Jackson, F. and Pargetter, R. (1986) "Ought, Options, and Actualism," *Philosophical Review* 95(2): 233-55.
- Louise, J. (2009) "I Won't Do It! Self-Prediction, Moral Obligation and Moral Deliberation," *Philosophical Studies* 146(3): 327-48.
- Mendola, J. (2006) *Goodness and Justice: A Consequentialist Moral Theory*, New York: Cambridge University Press.
- Portmore, D. W. (2013) "Perform Your Best Option," *The Journal of Philosophy* 110(8): 436-59.
- _____. (2011) *Commonsense Consequentialism*, Oxford: Oxford University Press.
- Ross, J. (2012) "Actualism, Possibilism, and Beyond," in M. Timmons, ed., *Oxford Studies in Normative Ethics, Vol. 2*, Oxford: Oxford University Press, pp. 74-97.
- Smith, H. (1978) "Doing the Best One Can," in A. Goldman and J. Kim, eds., *Values and Morals*, Dordrecht: Riedel, pp. 185-214.
- Sobel, J. H. (1976) "Utilitarianism and Past and Future Mistakes," *Noûs* 10(2): 195-219.
- Timmerman, T. (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.
- _____. and Cohen, Y. (2016) "Moral Obligations: Actualist, Possibilist, or Hybridist?" *Australasian Journal of Philosophy* 94(4): 672-86.
- Vessel, J. P. (2016) "Against Securitism, the New Breed of Actualism in Consequentialist Thought," *Utilitas* 28(2): 164-78.
- _____. (2009) "Defending a Possibilist Insight in Consequentialist Thought," *Philosophical Studies* 142(2): 183-95.
- Wedgwood, R. (2009) "Against Actualism," PEA Soup, peasoup.typepad.com/peasoup/2009/09/against-actualism.html (accessed Dec. 23, 2014).
- Woodard, C. (2009) "What's Wrong with Possibilism," *Analysis* 69(2): 219-26.
- Zimmerman, M. (2006a) "Is Moral Obligation Objective or Subjective?" *Utilitas* 18(4): 329-61.
- _____. (2006b) "The Relevant Risks to Wrongdoing," in K. McDaniel, J. Raibley, R. Feldman and M. Zimmerman, eds., *The Good, the Right, Life and Death: Essays in Honor of Fred Feldman*, Burlington, VT: Ashgate Publishing Co., pp. 151-72.

____ (1996) *The Concept of Moral Obligation*, Cambridge: Cambridge University Press.