

9

How to Be an Actualist and Blame People

Travis Timmerman and Philip Swenson

9.1 Introduction

The actualism/possibilism debate in ethics concerns the relationship between an agent's free actions and her moral obligations.¹ The core disagreement between these views is over whether an agent's free actions, in certain circumstances, at least partly determine her moral obligations. Actualists claim they do, while possibilists claim they do not. Although this debate intersects with questions of agency and moral responsibility (e.g., Are there true counterfactuals of freedom? Which of the agent's relevant act alternatives does blame track?), the more than forty-year-old literature contains shockingly little discussion of these issues. This is regrettable, in part, because the answers to questions about agency and responsibility bear directly on the plausibility of the views on offer in the actualism/possibilism debate and vice versa. This paper helps bridge that gap in the literature by assessing the plausibility of actualism and possibilism in light of *desiderata* about accounts of blameworthiness.

Our paper is structured as follows. In section 9.2 we provide a brief overview of the actualism/possibilism debate in ethics, defining each view and illustrating exactly how they come apart by considering a standard case from the literature. In section 9.3, we argue that standard forms of actualism are unable to straightforwardly accommodate a widely accepted, and exceedingly plausible, *desideratum* for accounts of blameworthiness. We refer to this as *Actualism's Blameworthiness Problem*. We then offer two solutions on behalf of the actualist in sections 9.4 and 9.5, showing how their view can, and should, be developed to avoid this problem. In section 9.6, we

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

show that actualism also has a *prima facie* problem providing a plausible account of blameworthiness for the outcomes of actions. We then offer a solution on behalf of the actualist. Finally, in section 9.7, we make a negative argument on behalf of the actualist by showing how possibilism is subject to its own, comparably troubling, blameworthiness problem. If our arguments succeed, they should collectively provide good reason to accept our suggested actualist strategies over alternative views.

9.2 An Overview of the Actualism/Possibilism Debate

It is uncontroversial among normative ethicists that the deontic status of an act depends upon the alternative acts available to the agent. As the actualism/possibilism debate illustrates, however, views about which acts are the *relevant* alternatives available to an agent are contentious. To get a grip on this abstract issue, consider the following first-order case:

The Gig: Brandi has been invited by her friend, Chad Kroeger, 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, Chad is a terrible musician. 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.² 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.³

² Purely for ease of presentation, we describe the actualist/possibilist cases in a way that implicitly assumes both counterfactual determinism and compatibilism. These assumptions are unnecessary for the purposes of our argument. Readers should feel free to amend the cases to be libertarian-friendly by replacing the “would” counterfactuals with “might” counterfactuals, where the agent is very likely to freely act in the way described. Those who reject counterfactual determinism can imagine that, instead of each of the agent's actions corresponding to a unique possible world, each of the agent's actions corresponds to a probability distribution of possible worlds that might be actualized if the agent performs the act in question.

³ This example is drawn from Cohen and Timmerman (2016: 1). It is structurally identical to standard *Professor Procrastinate*-cases from the literature. These types of cases appear throughout the literature, including Goldman (1978), Jackson and Pargetter (1986: 235), Vorobej (2000: 131–2), Portmore (2011: 180), and Timmerman (2015: 1512).

To be sure, Brandi *can* decide to attend the gig, and once there decide to be supportive of Chad. However, suppose that just isn't what Brandi *would* do if she were to attend. Here's the tricky philosophical question. Is Brandi morally obligated to accept or decline the invitation? *Actualists* hold that Brandi is obligated to decline because what would actually happen if she declines is better than what would actually happen if she accepts (Goldman 1976; Sobel 1976; Jackson and Pargetter 1986). *Possibilists* hold that Brandi is obligated to accept because accepting is part of the best possible sequence of acts that Brandi can perform (Greenspan 1978; Feldman 1986; Zimmerman 1996: §6; Zimmerman 2017).

While there is a variety of more precise technical formulations of these views on offer, we settle on what we take to be the most plausible versions of each view. This makes no difference to our arguments, however, since the issues we raise apply to the less precise, less plausible, definitions as well. We focus on two distinct formulations of actualism and one formulation of possibilism. First, consider a version of actualism that does not take into account which acts are under an agent's control at which times:

Actualism (without a control condition): At t an agent S is obligated to ϕ at t' (where t may or may not be identical to 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 perform any incompatible act-set S can perform at t .⁴

The earliest versions of actualism were formulated without a control condition,⁵ but have become much less popular since such versions were shown to incur a unique set of problems.⁶ To see how this version of actualism works, imagine that it is presently under Brandi's control to <accept the invitation and be supportive of Chad>. Nevertheless, it is true that Brandi wouldn't be supportive if she attends. A version of actualism without a control

⁴ As is standard in the literature, we formulate actualism and possibilism in terms of one's *objective*, rather than one's *subjective*, obligations. 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: 12–23)). To keep the dialectic as simple as possible, unless we state otherwise, we'll assume that the agents in the cases we give know all of the normatively relevant facts. This means that the agent's subjective and objective obligations are identical in such cases.

⁵ See, for instance, Goldman (1976), Sobel (1976), and Jackson and Pargetter (1986).

⁶ For discussions of seminal objections to actualism, see Wedgwood (2009), Baker (2012: 642–3), and Zimmerman (2017: §2). See also Cariani (2016: 409–10) and Timmerman and Cohen (2016: 674).

condition entails that Brandi is obligated to decline the invitation to the gig because what would happen if she declines is better than what would happen if she performs any incompatible act she can at the time in question (i.e., accept the invitation). This implication also illustrates one of the supposed problems with such versions of actualism. It is counterintuitive to hold that Brandi is obligated to decline the invitation if it's currently under her control to ensure that she does something much better, viz. both accept the invitation and be supportive. While such versions of actualism are admittedly subject to unique problems, as we demonstrate in section 9.3, they also allow for a unique benefit as well. Currently, the most popular forms of actualism include a control condition and focus on maximal acts. They may be formulated as follows. The underlined words highlight the differences between this formulation and the previous one:

Actualism (with a control condition): At t an agent S morally ought to φ at t' (where t may or may not be identical to t') iff φ -ing at t' is an act-set currently under S 's 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 incompatible maximal act-set under S 's control at t .^{7, 8}

To see how this version of actualism works, imagine again that it is presently under Brandi's control to <accept the invitation and be supportive of Chad>, but that it is not presently under her control to do anything more specific than that, including doing anything after the time at which she can support, or not support, Chad. As we're using the term, then, <accepting

⁷ Actualism without a control condition violates the principle of normative inheritance Portmore (2019: ch. 4) and it generates conflicting obligations without saying which obligation takes priority Cohen and Timmerman (2016: 11–12), Kiesewetter (2015: 929–34), and Portmore (2011: 181–3). Subsequent versions of actualism built in a control condition and focus on maximal act-sets to avoid this problem. See, for instance, Goldman (1978: 202), Carlson (1995: 121–3), Bykvist (2002: 61–4), and Jackson (2014). Douglas Portmore's (2011) and Jacob Ross's (2012) securitist views also count as versions of actualism for the purposes of this paper.

⁸ One might worry that we should not define actualism in a manner that assumes both consequentialism and that one ought to perform the best action available. (The same worry might apply to our account of possibilism below.) We make these assumptions, which are common in the literature on actualism, only for simplicity's sake. The intuitive thought behind actualism is that one should hold fixed what would result from one's action when deciding what one ought to do. But turning this thought into an account of right action is quite complicated without our simplifying assumptions. We think that actualists who reject these assumptions will nonetheless agree with our judgments about permissibility in the cases we appeal to below. So rejecting the simplifying assumptions should not undermine our argument.

the invitation and being supportive of Chad> is a maximal act-set because it is an act-set that is among the set of the most precise act-sets that Brandi can currently ensure she performs. More formally, an act-set α_i is maximal at time t iff there is no other set of actions α_j under the agent's control at t , such that performing α_i involves performing α_j , but not vice versa.⁹ Now, suppose that the maximal act-sets under S's control are <accept the invitation and be supportive of Chad>, <accept the invitation and *not* be supportive of Chad>, and <decline the invitation and not interact with Chad>. The non-maximal acts under Brandi's control are each of the conjuncts in the maximal act-sets. In a nutshell, when these are the maximal act-sets, actualism with a control condition entails that Brandi is obligated to <accept the invitation> because the best maximal act-set is <accept the invitation and be supportive of Chad> and bringing about that outcome requires that Brandi <accept the invitation>.

Actualism with a control condition will render a different verdict if it is not currently under Brandi's control to <accept the invitation and be supportive of Chad>. To illustrate, let's suppose that no matter what Brandi intends to do now, she would freely choose to not be supportive of Chad if she attends the gig. In this case, the maximal act-sets presently under Brandi's control are only <accept the invitation and *not* be supportive of Chad> and <decline the invitation and not interact with Chad>. Of these two alternatives, the best act-set is <decline the invitation and not interact with Chad>. Consequently, in this scenario, actualism with a control condition would entail that Brandi is obligated to <decline the invitation> because the best maximal act-set is <decline the invitation and not interact with Chad> and bringing about that outcome requires that Brandi <decline the invitation>. We will come back to the differences between the two standard versions of actualism in section 9.4. Finally, let's consider the standard form of possibilism:

Possibilism: At t an agent S is obligated to ϕ at t' (where t may or may not be identical to t') iff ϕ -ing at t' is part of the best series of acts that S can perform from t to the last moment that, at t , S can possibly perform an act.

⁹ This definition is borrowed from Portmore (2011: 177) and Portmore (2019). Jacob Ross employs this concept in his (2012), which he refers to as a "maximally specific option." There may be no maximally specific act-sets if, for any act-set an agent can perform, there is an infinite number of more precise ways they could perform that act-set. Although actualism and possibilism can easily be reformulated to accommodate this possibility (e.g., by discussing maximally specific *morally relevant* act-sets), we think doing so is unnecessary because actual agents lack the dexterity to perform infinitely precise variations of any act-set.

So, according to possibilism, Brandi is obligated to <accept the invitation> regardless of whether it is currently under her control to be supportive if she attends. This is simply because the best act-set she can, at t , perform is <accept the invitation and be supportive of Chad> and bringing about that outcome requires that Brandi <accept the invitation>. With these definitions in mind, we turn to the problem that blame presents for the actualist.

9.3 Actualism's Blameworthiness Problem

There are very few axioms about blame that are uncontroversial. Yet, the proposition that an agent is blameless if she has, relative to her evidence, done what was morally required of her for morality's own sake has a great deal of intuitive support. Many would agree that one is blameless for performing an act if (i) that act was subjectively morally permissible, (ii) the agent was motivated to perform the act for the right reasons (e.g., *because* it was the subjectively moral thing to do), and (iii) the agent was not motivated to perform the act for any of the wrong reasons (e.g., because it would harm an innocent person). Conditions (i)–(iii) appear to be jointly sufficient, but perhaps not necessary, conditions for merit-based blamelessness *for performing an act*.¹⁰ Given its extreme intuitive plausibility,¹¹ both normative ethical views and accounts of blameworthiness should want to accommodate some version of this *desideratum*, which we will refer to as the *Blameless Desideratum* (BD). BD may be formulated more precisely as follows:

Blameless Desideratum (BD): Agent S is blameless for φ -ing if φ -ing is itself subjectively permissible and if S φ s for all the right reasons and does not φ for any of the wrong reasons.¹²

¹⁰ We are not claiming, of any particular condition, that it is necessary for blamelessness. We also leave open that some subset of these conditions are sufficient for blamelessness. We also leave open the possibility that agents can deserve blame for something other than the performance of acts (e.g., having certain dispositions to act).

¹¹ Many philosophers hold that one can be blameworthy only for wrong acts. For a recent defense of this view see Montminy (2018). But even philosophers who hold that agents can be blameworthy for performing permissible acts say that such agents are blameworthy because they act for bad reasons or from bad motives (see Zimmerman (1997) and Capes (2012)). Thus BD is on solid footing. There is one motivation for rejecting BD in the literature: the claim that blameworthiness can *trace* back to previous bad choices. This will be discussed in section 9.5.

¹² BD is meant to be a plausible axiom for merit-based views of blame, not for consequentialist views of blame. There will certainly be cases where BD is satisfied, yet an agent ought to be blamed for consequentialist reasons and vice versa.

The problem for actualism is that standard forms appear unable to straightforwardly accommodate BD without committing themselves to deeply implausible consequences. To illustrate, consider the case of *Gamer Gabi*:

Gamer Gabi: Gabi has exactly \$500 in her account and initially planned to use it to purchase the latest and greatest tablet 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 tablet. It is not presently under Gabi's control to <keep the money in her account today and use the money to save the three lives tomorrow>. However, it is presently under her control to <keep the money in her account today>. Moreover, if she does keep the money in her account today, then it would be under her control tomorrow to <use the money to save the three lives>. At that point, all she has to do is to 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, suppose that Gabi is aware of these facts and consequently decides to purchase a tablet for herself today rather than tomorrow. She does so with the intention of avoiding any wrongdoing.¹³

Both versions of actualism suggest (assuming plausible claims about the extent of one's duties to aid others) that Gabi acts wrongly if she purchases a tablet tomorrow (rather than saving the three lives tomorrow), but acts permissibly if she purchases a tablet today (rather than saving the money to save the three lives tomorrow). So coupled with BD, both versions of actualism suggest that Gabi is blameless for purchasing a tablet today, but could be blameworthy for purchasing one tomorrow.¹⁴ This seems to be the

¹³ This example is drawn from Timmerman (2015: 1516–17) and Timmerman and Cohen (2016: 677). Those concerned about moral fetishism or schizophrenia may interpret Gabi's motivations to be too close to the motivations of the hospital visitor from Stocker (1976: 462) to fully satisfy BD. If Gabi is purely motivated by a *de dicto* desire to do good, then one may worry that she would not be acting for all the right reasons. See Smith (1994: 74–5) for a defense of this idea. In response, we will first note that we are unconvinced by these arguments against *de dicto* moral motivation, partly for reasons given in Shafer-Landau (1998) and Carbonell (2013). Providing a refutation of these arguments is beyond the scope of this paper, however. Second, while Gabi certainly is motivated by the good *de dicto*, she may be motivated by the good *de re* as well, at least within an actualist framework.

¹⁴ BD is formulated as a principle governing subjective permissibility and, as we have formulated it, actualism is a principle governing objective permissibility. But we think that any plausible formulation of subjective actualism will yield the result that, since she is aware of the relevant facts, Gabi acts *subjectively* wrongly if she purchases a tablet tomorrow, but acts *subjectively* permissibly if she purchases a tablet today. So the problem will still emerge. The same goes for the *Torturing Tammy* case discussed below.

incorrect result. At a minimum, Gabi seems to be equally blameworthy (or blameless) whether she purchases a tablet today or tomorrow. Moreover, it is intuitive to hold that Gabi is deserving of blame if she purchases a tablet today or tomorrow because it remains true of Gabi, both today and tomorrow, that she *can* abstain from acquiring a trivial good for herself in order to save three lives.¹⁵

Actualism's problem with accommodating BD stems from the fact that, according to actualism, an agent's *blameworthy behavior* in counterfactual situations is explanatorily prior to some of the agent's *moral obligations*. Stated oversimplistically, according to actualism, the fact that an agent would perform a blameworthy immoral act in a given situation often allows the agent to avoid incurring an obligation to do the best she can in said situation. In *Gamer Gabi*, the fact that Gabi would culpably act immorally when faced with the opportunity to spend \$500 to save three lives allows her to avoid incurring an obligation to use her \$500 to save the three lives. Specifically, the fact that she would culpably act wrongly if she were in that situation factors into the evaluative ranking of purchasing a tablet today and, since Gabi would not save the three lives tomorrow if she has the opportunity, actualism entails that Gabi is not obligated to keep the money in her account today.

If it were true, however, that Gabi would save the three lives tomorrow were she to have the \$500, then actualism would entail that Gabi is obligated to keep the money in her account today and use it tomorrow to save the three lives. So, Gabi's blameworthy immoral behavior of failing to save the three lives is explanatorily prior to her obligations to keep (or not keep) the money in her account today. This already seems to many to be implausible. But it is even more implausible if the agent's dispositions to do wrong *also* allows them to avoid blame. This problem can be most clearly be seen by considering a variation of a case from Zimmerman's (2017: 121):

Torturing Tammy: Tammy likes to torture. She's just captured a stray cat and is wondering what she ought to do with him. She can easily let him go and refrain from torturing any animal from here on out. If she lets him go, however, she would later choose to capture him again and torture

¹⁵ This intuition seems to us to be widely held among actualists and non-actualists alike. Anyone who lacks this intuition, however, poses no problem for the actualist since those who lack it will not think actualism has a blameworthiness problem.

him for ten minutes. If she instead tortures him for five minutes now, she will choose to not capture him or torture him again. Knowing this and wanting to avoid any wrongdoing, Tammy tortures the cat for five minutes.

According to both types of actualism (when combined with any plausible normative ethical theory), Tammy is obligated to torture the cat for five minutes, even though (the actualist agrees) she can simply not torture the cat at all. The fact that Tammy would culpably act wrongly if she doesn't now torture the cat is explanatorily prior to her incurring an obligation to torture the cat for five minutes. So, assuming BD, Tammy may then be blameless for torturing the cat for five minutes even though she could have simply not tortured it at all.

There may be an additional worry for the actualist here. Suppose Tammy's act of torturing the cat now is highly self-sacrificial. She would be horrified by the experience of torturing the cat now, but, in a different mood, would highly enjoy it later. She decides to have the horrifying experience now in order to minimize the total suffering of the cat. In this scenario, there may be pressure on the actualist to say that Tammy is *praiseworthy* for torturing the cat now. After all, she did the right thing for the right reasons, at great personal cost to herself. This seems like an even worse result.¹⁶

To recap, *Torturing Tammy* and *Gamer Gabi* are instances of *Actualism's Blameworthiness Problem*. They show that actualism coupled with BD seemingly generates extremely counterintuitive results. Actualism gets these results because it makes agents' potential blameworthy behavior explanatorily prior to their moral obligations. The most straightforward ways to avoid these problems is by either rejecting actualism or by rejecting BD outright. Our goal, however, is to do neither. Rather, we offer the best defenses available to the actualist that, in spite of initial appearances, allows them to accommodate BD (or something very close to it) and keep the essential components of their view intact. In sections 9.4 and 9.5, we will offer two defenses on behalf of the actualist.

¹⁶ Thanks to Justin Coates for pointing out this possibility to us. A referee suggested the possibility that perhaps Tammy is not praiseworthy full stop, but only "praiseworthy in comparison" (perhaps this just means less blameworthy) to what her moral status would have been had she instead performed the worse act of torturing later. The worry, as we see it, is that the actualist doesn't have a good explanation of why Tammy isn't praiseworthy full stop. She engages in significant self-sacrifice in order to act rightly for the right reasons. What more is required for praiseworthiness? But if the referee's suggestion is correct the praiseworthiness worry is avoided. Since our goal is to defend actualism, we would be happy with this result.

9.4 First Solution: Adopt a Contextualist Form of Actualism

One way actualism can accommodate BD without generating counterintuitive consequences is by finding a way to identify (subjective) wrongdoing in *Gamer Gabi* and *Torturing Tammy* cases within the actualist framework. This may seem impossible since actualism affirms that Tammy is obligated to torture the cat for five minutes and denies that Gabi is obligated to keep the money in her account to save the three lives tomorrow. Despite initial appearances, however, this task is indeed possible. The actualist can identify (subjective) wrongdoing within an actualist framework by positing a plurality of actualist “oughts,” where at least one of those oughts entails that agents act wrongly in *Gamer Gabi* and *Torturing Tammy*-style cases. We will refer to such views as *contextualist actualism*. The basic idea is that, in addition to standard salient actualist obligations, agents have actualist obligations to perform the best act-set they can over the course of their life. We will refer to this actualist obligation as a *lifetime actualist obligation*. Agents will necessarily violate this lifetime actualist obligation in any case structurally identical to *Gamer Gabi* or *Torturing Tammy*. Frank Jackson and Robert Pargetter (1986) developed and defended the best version of contextualist actualism, which may be formulated as follows:¹⁷

Contextualist Actualism (CA): whether *S* ought to φ at *t* depends on whether φ -ing at *t* is better than what *S* would do if *S* does not φ at *t*. But what *S* ought to do at *t* is whichever of the maximally relevantly specific options available to *S* at *t* would have the best outcome.¹⁸

To better understand how CA identifies (subjective) wrongdoing, consider its assessment of *Torturing Tammy* and, to keep the case as simple as possible, assume that Tammy will suffer an inevitable death immediately after the time at which she can torture or not torture animals. Whether Tammy ought to torture the cat for five minutes depends on whether torturing the cat for five minutes is better than what Tammy would do if she were to not torture the cat for five minutes. Since Tammy would torture the cat for ten minutes if she doesn't torture it for five, it follows that Tammy

¹⁷ For another version of contextualist actualism, see McKinsey's (1979).

¹⁸ Our formulation of Jackson and Pargetter's position is borrowed from Ross's (2013: 76).

ought to torture the cat for five minutes. This commitment of CA is simply a salient standard actualist verdict.

Now, consider the second part of the definition. *What* Tammy ought to do is completely refrain from torturing the cat since that is the maximally relevantly specific option available to her that would have the best outcome. That is simply the best thing that Tammy can do over the course of her life. To be sure, as Jackson and Pargetter were conceiving of CA, *what* an agent ought to do is co-extensive with the possibilist obligation.¹⁹ Nevertheless, this ought does admit of an actualist justification. After all, what would *actually* happen if Tammy were to <completely refrain from torturing the cat> is better than what would actually happen if Tammy were <not to completely refrain from torturing the cat> or to do anything else at the times in question. Thus, *what* Tammy ought to do is not torture the cat at all and she ought to do this for actualist reasons.²⁰

CA renders a similar verdict in *Gamer Gabi* cases. *Whether* Gabi ought to keep the money in her account today depends on whether doing so would result in a better outcome than spending the money on a tablet today. It wouldn't since Gabi wouldn't use the money to save the three lives regardless of her choice today. So, we can suppose that each outcome is equally good, which means that the *whether*-ought doesn't favor either action. *What* Gabi ought to do, however, is to <keep the money in her account today and use the money to save the three lives tomorrow>. By purchasing a tablet today, Gabi violates her (subjective) lifetime actualist obligation and so fails to meet BD's sufficient conditions for blamelessness. Gabi may then be blameworthy for purchasing a tablet today and just as blameworthy for purchasing one today as for purchasing one tomorrow. This is the correct

¹⁹ Those who prefer a version of actualism with a control condition will hold that the maximally relevantly specific option available to S is identical to this actualist obligation, rather than the possibilist obligation. Such actualists do not regard the possibilist obligation as a relevant option for the agent when it is not currently under the agent's control to ensure that she fulfills the possibilist obligation.

²⁰ The careful reader will notice that actualism (without a control condition), as formulated, is consistent with CA. It even entails something close to CA since it contains no restrictions on which act-sets can stand in for " φ " or what times can stand in for " t ". To illustrate, in *Torturing Tammy*, when φ denotes <torturing the cat for five minutes> and t denotes the five minutes Tammy would torture the cat, actualism without a control condition entails that Tammy ought to <torture the cat for five minutes>. However, when φ denotes <completely refraining from torturing the cat> and t refers to the remaining time in her life, actualism without a control condition entails that Tammy ought to <completely refrain from torturing the cat>. On the other hand, actualism (with a control condition) is inconsistent with CA. Since it compares the goodness of the act in question to alternative *maximal* act-sets, it will never generate conflicting actualist oughts.

result. After all, whether she purchases a tablet today, or tomorrow, she will violate the same moral obligation (i.e., her lifetime actualist obligation) and will bring about an equally bad outcome (i.e., three innocent deaths).

As these cases illustrate, adopting CA allows the actualist to accommodate BD without generating the incorrect verdicts in *Gamer Gabi* and *Torturing Tammy* cases. This is an overlooked advantage of CA. But it may not be enough to attract many actualist converts for a few reasons. First, if blame-worthiness tracks the lifetime actualist obligation, then CA's account of blame may be identical to the possibilist's account of blame. As we will demonstrate in section 9.7, there are plausible reasons to reject that blame tracks the lifetime actualist obligation, reasons that actualists are likely disposed to accept.

Second, CA fails to be action-guiding because it generates conflicting moral oughts and furthermore is silent about which ought takes priority. Although CA holds that Tammy ought to <torture the cat for five minutes>, CA also holds that Tammy ought to <refrain from torturing the cat altogether>. These incompatible actions prescribed by CA seem to put Tammy in an action dilemma (and perhaps a blame dilemma). Jackson recognized this worry in his (2014), but argued that there is no action dilemma because, at any given time, agents can only act in accordance with the *whether*-ought. They cannot ensure that they act in accordance with the *what*-ought when it concerns actions not under the agent's control at the time in question (2014: 636). The takeaway is that Jackson identifies the *whether*-ought as the action-guiding one.

Jackson's response is unsatisfactory for two reasons. First, designating the *whether*-ought as the action-guiding one entails that the agent will violate her *what*-ought prescriptions whenever they conflict with her *whether*-ought prescriptions. Most actualists will want to deny that agents have action-guiding moral obligations that require them to perform wrong actions. Second, notice that CA also generates conflicting obligations when applied to synchronic and diachronic act-sets, where performing each possible act-set is under the agent's control at the time in question. In such cases, the agent will be able to act in accordance with either mutually exclusive ought at the time in question, yet (as formulated) CA still doesn't say which one takes priority. Following Jackson's (2014), the reader might infer that the *whether*-ought takes priority in such cases. But that would be deeply counterintuitive because acting in accordance with the *what*-ought guarantees that the agent acts in accordance with each of her moral oughts, while acting in accordance with the *whether*-ought in such cases ensures that

one violates the *what*-ought prescriptions. Thus, in cases where it is presently under the agent's control to fulfill both the *whether*-ought and the *what*-ought prescriptions, it seems that the *what*-ought should take priority. This is because, in those cases, fulfilling either ought will presently be under the agent's control and only fulfilling the *what*-ought guarantees that the agent acts in accordance with each of the moral oughts.

Of course, proponents of CA could designate the *whether*-ought as action-guiding when it's not presently under the agent's control to fulfill the *what*-ought prescription and designate the *what*-ought as action-guiding in whenever it is presently under the agent's control to act in accordance with it. Though seemingly ad hoc, it may be a promising strategy for CA, but one that we don't expect to be dialectically effective against actualists who prefer a formulation with a control condition. In light of these considerations, in section 9.5, we develop another way for such actualists to accommodate BD without generating intuitively incorrect verdicts in *Gamer Gabi* and *Torturing Tammy* cases.

9.5 Second Solution: Appeal to Tracing

A second option for the actualist is to reject BD as formulated, but to sugar the pill by positing an alternative, not-implausible, principle in its place. We think this is a promising strategy because philosophers have provided independent motivation for an alternative position, appealing to reasons that are entirely independent of actualism. Specifically, many philosophers accept the idea that one can be blameworthy for an act that can be *traced* back to a previous blameworthy choice.²¹ To illustrate, consider this case:

Self-Deluded Daisy: Daisy wants to skip her distant relative's birthday party on Saturday to instead play golf. She previously promised to go to the party, but at this point is no longer interested in attending. Still, Daisy does not want to be blameworthy for missing the party. So, she takes a drug on Friday that causes her to forget that she promised to attend. The next day she skips the party and plays golf with a clear conscience.

It is relatively uncontroversial that Daisy is blameworthy for <taking the drug>. Some philosophers, however, also believe that Daisy is blameworthy

²¹ For an excellent recent defense of tracing see Fischer and Tognazzini (2009).

for <skipping the party>. This verdict appears to conflict with BD. This is because, given her memory loss, Daisy's <skipping the party> is subjectively permissible and we'll stipulate that, on Saturday, Daisy acts only for the right reasons. Nevertheless, many find it intuitive that Daisy is blameworthy for <skipping the party> in virtue of her choice to take the drug yesterday. By accepting this claim, one is accepting that Daisy's blameworthiness is *traced* back to her previous bad choice. Those who agree with this judgment must grant that BD is false and hold that, in certain cases, agents can be blameworthy for performing a subjectively permissible act for all of the right reasons and none of the wrong reasons.

To better understand how the tracing strategy can help the actualist, it will be useful to distinguish between direct and indirect blameworthiness, which we define as follows:

Direct Blameworthiness: Agent *S* is directly blameworthy for φ -ing if she is (i) blameworthy for φ -ing and (ii) the blameworthiness for φ -ing does not trace back to previous acts.

Indirect Blameworthiness: Agent *S* is indirectly blameworthy for φ -ing if she is (i) blameworthy for φ -ing and (ii) the blameworthiness for φ -ing does trace back to previous acts.²²

With this distinction in hand, we can formulate a version of BD that is consistent with *tracing*:

*BD**: Agent *S* is *directly* blameless for φ -ing if φ -ing is itself subjectively permissible and if *S* φ s for all the right reasons and does not φ for any of the wrong reasons.

Agent *S* is *indirectly* blameless for φ -ing iff there are no previous acts which *S* is directly blameworthy for performing to which φ -ing traces back.

*BD**'s implications for *Torturing Tammy* and *Gamer Gabi* depends on how we fill in the details. Specifically, it depends on whether Tammy and Gabi are culpable for certain facts. To illustrate, consider two ways of filling in Gabi's backstory:

²² This distinction is common in the literature. See, for example, Rosen (2004: 298–300) and Clarke (2017: 235).

Traceable Gabi: Traceable Gabi is to blame for the fact that she would purchase the tablet tomorrow. She would purchase the tablet tomorrow because she has bad dispositions for which she is morally responsible for possessing. Her bad dispositions are the result of her repeatedly making selfish choices in the past.

Untraceable Gabi: Untraceable Gabi is not to blame for the fact that she would purchase the tablet tomorrow.

The actualist who replaces BD with BD* has the resources to say that *Traceable Gabi* is blameworthy for purchasing a tablet today because *Traceable Gabi* is indirectly blameworthy for purchasing the tablet despite the fact that doing so is subjectively permissible.²³ To be sure, the actualist will still be unable to hold that *Untraceable Gabi* is blameworthy. However, the actualist should insist that the much of the motivation for judging that Gabi is blameworthy is due to the assumption that Gabi is responsible for the fact that she would buy a tablet tomorrow. Once we drop that assumption, it becomes more plausible that Gabi is not to blame.

Many will resonate with BD*'s verdict in *Untraceable Gabi* because it takes into account the fact that, unlike *Traceable Gabi*, *Untraceable Gabi* couldn't have done anything to develop a better moral character than the one she, in fact, has. Although this is controversial, there is something attractive about the thought that Gabi is not blameworthy for minimizing the damage wrought by a moral character that she was herself powerless to improve.²⁴ We grant that there is some lingering counterintuitiveness to the claim that *Untraceable Gabi* is not blameworthy for <purchasing a tablet today>. But it is much less implausible than the claim that *Traceable Gabi* is

²³ Philosophers who reject tracing might still appeal to Traceable Gabi's past selfish choices to explain away the intuition that she is blameworthy for purchasing the tablet. The fact that Traceable Gabi's selfish choices would lead to her current circumstances might suggest that she is more blameworthy for those past choices than she otherwise would have been. This explains why we take her current behavior to be relevant to Gabi's degree of blameworthiness (even though she is not, strictly speaking, blameworthy for purchasing the tablet). Thanks to David Brink and Andrew Forcehimes for suggesting this sort of approach.

²⁴ Some philosophers hold that agents can be directly blameworthy for possessing traits which they are powerless to avoid possessing (see e.g., Smith (2005)). On this view it is plausible that Untraceable Gabi does not get off the hook for purchasing the tablet because she cannot avoid possessing her bad character. For further discussion of this sort of issue, see Shoemaker (2015: chs 1 and 5). We are inclined to reject the claim that agents can be directly responsible for possessing traits which they are powerless to avoid. But defending our position here would be beyond the scope of this paper.

not blameworthy. Thus, actualists who accept tracing and reject BD in favor of BD* have, in our view, made considerable progress.

9.6 Actualism and Blameworthiness for Outcomes

Actualism also faces a *prima facie* difficulty with respect to blameworthiness for outcomes (as opposed to actions). To illustrate, consider the following case:

Self-Undermining Rescue: John is walking along the beach and sees a child drowning in the water. John believes that he could rescue the child without much effort. Due to his laziness, he decides not to attempt to rescue the child. The child drowns. Unbeknownst to John, if he had jumped into the water and begun swimming toward the child, it would have occurred to him that he was late for an appointment. John would have then freely decided to turn back to shore and he would not have rescued the child.

In *Self-Undermining Rescue* it seems that John is blameworthy for the death of the child. Contrast that with cases that have the following structure:

Sharks: John is walking along the beach and sees a child drowning in the water. John believes that he could rescue the child without much effort. Due to his laziness, he decides not to attempt to rescue the child. The child drowns. Unbeknownst to John, there is a school of sharks hidden beneath the water. If John had jumped in and attempted to rescue the child, the sharks would have killed him and his rescue attempt would have been unsuccessful.²⁵

In *Sharks* it seems that John is *not* blameworthy for the death of the child. (Though, of course, he is blameworthy for not trying to save the child.) Our differing judgments in *Self-Undermining Rescue* and *Sharks* pose two related problems for actualism.

First, actualism seems to imply that John's refraining from jumping in is (objectively) morally permissible in both cases. After all, things would not have turned out any better if John had jumped in and, we can stipulate, that

²⁵ This case is drawn from Fischer and Ravizza (1998: 125).

it was not presently under John's control to <jump in and save the child> in either case. Here's the problem. If John's omission is (objectively) permissible in *Self-Undermining Rescue*, then it is hard to see why he is blameworthy *for the death of the child*.

Second, unlike the sharks in *Sharks*, it seems unintuitive to treat John's hypothetical future choice as an obstacle outside of his control. After all, John seems blameworthy for the death of the child in this case. Obstacles outside of one's control should be held fixed when determining one's blameworthiness (e.g., just as the presence of the sharks are held fixed when determining whether John is blameworthy for the death of the child in *Sharks*.) Standard judgments in *Self-Undermining Rescue* and *Sharks* seem to fit better with the way possibilism treats future choices than it does with the way actualism treats future choices. In other words, possibilism can straightforwardly account for the asymmetric judgments in these two cases, while actualism cannot.

How should actualists respond to these worries? They should begin by pointing out that, although John's omission was objectively permissible, his omission was subjectively wrong in both cases. Again, subjective obligations pick out what a conscientious moral agent should do, relative to the evidence in her epistemic ken. Given that John believed that he would rescue the child if he jumped in the water, he was subjectively obligated to jump in the water. So, John acts subjectively wrongly in both cases. This explains why John is blameworthy for not jumping in and attempting a rescue in both cases. But what explains why he is intuitively blameworthy *for the death of the child* in *Self-Undermining Rescue* but not in *Sharks*?

One option for the actualist is to again appeal to tracing. Perhaps the intuition that John is blameworthy for the death in *Self-Undermining Rescue* is based on the assumption that John is responsible for being the sort of person who would choose to allow a child to drown in order to be on time for an appointment. That is, John is responsible for the fact that his attempt to rescue the child in *Self-Undermining Rescue* would not succeed. In *Sharks*, however, John is not responsible for the fact that his attempt to rescue the child would not succeed. This may explain our differing intuitive judgments. If so, actualists can appeal to this difference between the cases to keep their view intact and avoid such counterintuitive consequences.

Even with this promising strategy, the actualist isn't completely immune from this sort of objection. It is important to recognize that the tracing approach does not account for John's responsibility for the death of the child in *Self-Undermining Rescue* if we stipulate that John is not responsible for

the fact that he would have turned back. But, as with *Gamer Gabi*, we think that the actualist should claim that the reasons to think John is blameworthy for the death of the child are much stronger in the version of the case that allows for tracing. They could simply bite the (relatively small) bullet in the remaining cases.

Some actualists may not want to bite this bullet, however, and they needn't do so. Actualists can account for John's blameworthiness in versions of *Self-Undermining Rescue* where he is not responsible for the fact that he would have turned back in the following way. They can hold that John's responsibility for the death of the child is explained by the fact that he *could have* performed a sequence of acts, each of which he would have been subjectively obligated to perform (had he performed each of the previous acts in the sequence) that would have resulted in the survival of the child. In *Sharks*, however, there is no sequence of acts that John could have performed (subjectively required or otherwise) that would have resulted in the survival of the child. More precisely, the actualist can endorse this principle:

Blame For Outcomes: S is blameworthy for outcome O if (1) O occurs and is bad. (2) S meets the epistemic requirement for responsibility with regard to O. And (3) there is a sequence of acts such that (a) S was subjectively obligated to perform the first act in the sequence, (b) S would have been subjectively obligated to perform each act in the sequence if S had performed all previous acts in the sequence, and (c) performing all acts in this sequence would have prevented O.

Blame For Outcomes is perfectly consistent with actualism and explains why John is blameworthy for the death of the child in *Self-Undermining Rescue*, but not in *Sharks*. It also seems to us to be independently plausible. For these reasons, we believe that the actualist should endorse it.

One might object that *Blame For Outcomes* is possibilist in spirit and that it is thus ad hoc for the actualist to endorse it. We want to resist this objection. The dispute between actualism and possibilism is so difficult, in part, because both views seem to be getting at something important about morality. So, each person should want the view they endorse to accommodate, so far as is possible, intuitive insights associated with the other view. The actualist can accommodate the intuitive judgment about *Self-Undermining Rescue* by endorsing *Blame For Outcomes* for principled reasons and we think they should do so. Although we think this is the route actualists should take, as we already mentioned, it remains open to

the actualist to simply opt for the first strategy we outline and bite the bullet in the cases in which John is not responsible for the fact that he would have turned back.

9.7 Possibilism's Blameworthiness Problem

The best defense is a good offense. We will now go on the offensive on behalf of the actualist. The worry we raised for actualism is that it lets agents like Gabi and Tammy off the hook too easily. One might think that possibilism demands more from agents and thus will not be subject to this sort of worry. But, as we will now argue, possibilism also seems to let agents off the hook too easily. Possibilism can be precisified to avoid letting agents off the hook too easily, but doing so causes possibilism to incur other non-trivial difficulties. Consider the case of *Lazy Larry*:

Lazy Larry: Larry believes that he will face a choice tomorrow between donating \$500 to save three lives or spending the money on himself. He believes that it is very important that he donate the \$500. But he also believes that he will only donate the \$500 if he spends the evening reading essays by Peter Singer. Otherwise, he will act on his desire to spend the money on himself. He believes that he could freely donate the money even if he does not read Singer, but he is confident that he would freely choose not to donate if he doesn't read Singer. Unfortunately, Larry hates reading philosophy, but he loves playing video games. He also knows that playing video games will give his opponents a tiny bit of extra joy (one hedon), so he decides to spend the evening playing video games because he wants to do the morally right thing.²⁶

It seems quite obvious that Larry is blameworthy for not reading the Singer essays and he seems blameworthy regardless of whether he ends up donating the \$500 to save the three lives. However, it is not easy for the possibilist to account for this. Given how displeasing Larry finds philosophy and the

²⁶ The worry we raise based on *Lazy Larry* is related to a potential problem for possibilism which Smith (2014: 20) suggests in a footnote. She writes that possibilism "seems ill equipped to generate appropriate duties to gather information. Since . . . many . . . agents faced with a decision whether or not to acquire more information [are] able—whether or not [they] investigate . . . —to do the right thing in the future, a possibilistic version of the duty to investigate would not generate any objective duty for [them] to gather evidence."

benefit that results from his playing video games, it looks as though the best act-set Larry could perform is <spend the evening playing video games and donate the \$500 tomorrow>.

Thus according to possibilism, Larry is (subjectively and objectively) obligated to refrain from reading philosophy. So, it is hard to see how he could be blameworthy for refraining. One initial move the possibilist might make would be to agree that, according to possibilism, Larry is not to blame for not reading the Singer essays, but claim that this is not a problem for possibilism because Larry will be highly blameworthy soon enough. Once tomorrow rolls around and Larry fails to donate, he will be blameworthy. Thus, the possibilist might claim, possibilism does not let Larry off the hook too easily.

We think that there is still a cost in denying that Larry is blameworthy for <not reading the Singer essays> even if he is still blameworthy for <failing to donate \$500 tomorrow>. But there is a deeper problem for the possibilist here. Suppose that Larry's beliefs about tomorrow are mistaken. Unbeknownst to him, when he is faced with the choice tomorrow, he will muster up enough willpower to avoid succumbing to *akrasia* and donate the \$500 to save the three lives. Given this, possibilists and actualists will agree that Larry is *objectively* permitted to spend the evening playing video games. They will only disagree about whether it is *subjectively* permissible to do so. Possibilists will hold that Larry is subjectively obligated to play video games, while actualists will hold that Larry is subjectively obligated to read Singer.

As noted earlier, we initially formulated possibilism in terms of objective obligations. However, it can be formulated in terms of subjective obligations too. Here is one such formulation:

Subjective Possibilism: At t an agent S is subjectively obligated to φ at t' iff (i) S believes she can φ at t' and (ii) S believes that φ -ing at t' is part of the best series of acts that S can perform from t to the last moment that, at t , S can possibly perform an act.²⁷

On subjective possibilism, Larry is subjectively permitted to refrain from reading Singer even though he (falsely) believes that he would save the three

²⁷ Some philosophers prefer to formulate subjective obligations in terms of what S 's evidence supports or what S "should believe" rather than in terms of what S actually believes. Our objection to possibilism would not be substantially affected by switching to an evidential support account of subjective obligation.

lives if and only if he reads Singer. Thus, it looks as though BD entails that he is blameless, and this is the incorrect result. Note that the possibilist cannot take solace in the claim that Larry will eventually be blameworthy for failing to donate since Larry won't necessarily fail to donate the money.

Perhaps the possibilist should insist that Larry *is* to blame for not reading Singer even though doing so was permissible. This would allow the possibilist to avoid the charge of letting Larry off the hook too easily, but it would require them to reject BD. (They would have to reject BD* as well as since Larry's situation need not trace back to any previous blameworthy acts on his part.) On this approach, Larry is blameworthy for playing video games, not because playing video games is itself impermissible, but because performing this permissible act would seemingly result in Larry performing a deeply impermissible act in the future, one that is worse than the trivially impermissible act he would seemingly have to perform to avoid the future wrongdoing. The possibilist could amend BD as follows to accommodate this judgment. The underlined words highlight the differences between this formulation of BD and the original formulation:

*BD***: Agent *S* is blameless for φ -ing if φ -ing is itself subjectively permissible and will (relative to the agent's evidence) only result in *S* performing subjectively permissible acts and if *S* φ s for all the right reasons and does not φ for any of the wrong reasons.

We think rejecting BD in favor of BD** (and saying Larry is blameworthy for his permissible act) may be the best recourse for the possibilist, but it does raise serious difficulties for the view. First, it threatens to commit the possibilist to the existence of blame dilemmas. Suppose that Larry ends up choosing to read Singer. Should the possibilist say that Larry is also blameworthy for reading Singer? If they say yes, then they commit themselves to the existence of blame dilemmas. Larry is blameworthy whether or not he reads Singer. This may not seem so counterintuitive if Larry is in a *self-imposed* blame dilemma (e.g., being required to fulfill two incompatible promises freely made). But that needn't be the case, as Larry need not be at fault for being in his current situation. This version of possibilism could allow agents to be in *world-imposed* blame dilemmas.

Suppose instead that the possibilist says that Larry is not to blame if he reads Singer. They are then endorsing a very odd claim, viz. Larry is blameworthy if he does what is subjectively permissible (playing video

games) and blameless if he does what is subjectively wrong (reading Singer). More generally, they commit themselves to the following:

Blameworthy Right Action & Blameless Wrong Action: An agent *S* can be in a circumstance where *S* would be blameworthy for fulfilling their subjective moral obligation (for the right reasons) and blameless for freely violating their subjective obligations.

We find this consequence quite counterintuitive. It is especially strange because Larry does not seem to have an excuse for doing what is subjectively wrong (i.e., reading Singer). So why isn't he blameworthy for performing the subjectively wrong action?

So, we see three viable options for the possibilist: (1) bite the bullet and deny that Larry is blameworthy for playing video games, (2) accept world-imposed blame dilemmas, (3) accept *Blameworthy Right Action & Blameless Wrong Action*. Each of these have obvious costs.

In our view, the best option for the possibilist is (2). They should replace BD with BD** and commit themselves to world-imposed blame dilemmas. While this is the best (or least bad) strategy for the possibilist, it still comes at a non-trivial cost for the view, one which we think illustrates that possibilism is subject to its own blameworthiness problem on a par with actualism's blameworthiness problem.

9.8 Conclusion

We sought to establish a few related theses in this paper. We first illustrated that existing forms of actualism, coupled with BD, are unable to straightforwardly account for why agents such as *Gamer Gabi* are blameworthy for their actions. We then offered two independent solutions on behalf of the actualist, each with their own advantages and limits. First, one could adopt an explicitly contextualist form of actualism, thereby allowing actualists to hold that agents such as Gabi do violate BD. The second option for the actualist is to adopt an account of tracing that will allow them to account for such agent's blameworthiness in cases in which these agents are responsible for their rotten moral dispositions.

In section 9.6, we demonstrated that actualism also has difficulty straightforwardly accommodating plausible judgments about blameworthiness for outcomes. Specifically, it has trouble accommodating the

asymmetrical judgments between *Self-Undermining Rescue* and *Sharks*. We again offered two solutions on behalf of the actualist. First, we applied the tracing strategy outlined in section 9.5 to show how it offers a parallel solution to this problem. Second, we demonstrated that adopting our *Blame for Outcomes* principle allows actualists to accommodate these asymmetrical judgements.

In section 9.7, we went on the offensive, and argued that possibilism, coupled with BD, has trouble straightforwardly accounting for why agents such as *Lazy Larry* are blameworthy for their actions. In light of our case, the possibilist will have to either (1) bite the bullet (2) commit themselves to world-imposed blame dilemmas or (3) accept *Blameworthy Right Action & Blameless Wrong Action*. We suggested that (2) is the best recourse for the possibilist. Our arguments should collectively provide reason to believe that, contrary to initial appearances, actualism may be able to adequately accommodate the highly plausible blameworthiness *desiderata* discussed in this paper.²⁸

Works Cited

- Baker, Derek. (2012). "Knowing Yourself—And Giving up on your Own Agency in the Process." *Australasian Journal of Philosophy* 90(4): 641–56.
- Bykvist, Krister. (2002). "Alternative Actions and the Spirit of Consequentialism." *Philosophical Studies* 107(1): 45–68.
- Capes, Justin A. (2012). "Blameworthiness without Wrongdoing." *Pacific Philosophical Quarterly* 93(3): 417–37.
- Carbonell, Vanessa. (2013). "De Dicto Desires and Morality as Fetish." *Philosophical Studies* 163(2): 459–77.
- Cariani, Fabrizio. (2016). "Consequence and Contrast in Deontic Semantics." *Journal of Philosophy* 96(3): 396–416.
- Carlson, Erik. (1995). *Consequentialism Reconsidered* (Dordrecht: Riedel).

²⁸ For very helpful feedback on earlier drafts of this paper, we wish to thank Sean Aas, Marius Baumann, Gunnar Björnsson, David Brink, Justin Coates, Julia Driver, Andrew Forcehimes, Camil Golub, Elizabeth Harman, Matthew Hammerton, Robert MacDougall, August Gorman, Thaddeus Metz, Collin O'Neil, Doug Portmore, Luke Semrau, David Shoemaker, Michael Smith, Stephen White, two anonymous referees, the audience at the Neutrality: Reasons, Values, and Times conference at Nanyang Technological University in Singapore, and the audience at the New Orleans Workshop on Agency and Responsibility (NOWAR). Work on this paper was funded by summer grants from The College of William & Mary and Seton Hall University, as well as a grant from the National Endowment for the Humanities.

- Clarke, Randolph. (2017). "Ignorance, Revision, and Commonsense." In P. Robichaud and J. W. Wieland (eds), *Responsibility: The Epistemic Condition* (New York: Oxford University Press) pp. 233–51.
- Cohen, Yishai and Timmerman, T. (2016). "Actualism Has Control Issues." *The Journal of Ethics and Social Philosophy* 10: 1–18.
- Feldman, Fred. (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic* (Dordrecht: D. Reidel).
- Fischer, J. M., and Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge: Cambridge University Press).
- Fischer, John Martin and Tognazzini, Neal A. (2009). "The Truth about Tracing." *Noûs* 43(3): 531–56.
- 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 A. I. Goldman and J. Kim (eds), *Values and Morals* (Dordrecht: Dordrecht) pp. 185–214.
- Greenspan, P. S. (1978). "Oughts and Determinism: A Response to Goldman." *The Philosophical Review* 87(1): 77–83.
- Jackson, Frank. (2014). "Procrastinate Revisited." *Pacific Philosophical Quarterly* 95(4): 634–47.
- Jackson, Frank and Robert Pargetter. (1986). "Ought, Options, and Actualism." *The Philosophical Review* 95(2): 233–55.
- Kiesewetter, Benjamin. (1979). "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125(4): 921–46.
- McKinsey, Michael. (1979). "Levels of Obligation." *Philosophical Studies* 35(4): 385–95.
- Montminy, Martin. (2018). "Culpability and Irresponsibility." *Criminal Law and Philosophy*: 167–81.
- Portmore, Douglas W. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality* (New York: Oxford University Press).
- Portmore, Douglas W. (2019). *Opting for the Best: Oughts and Options* (New York: Oxford University Press).
- Rosen, Gideon. (2004). "Skepticism about Moral Responsibility." *Philosophical Perspectives* 18: 295–313.
- Ross, Jacob. (2013). "Actualism, Possibilism, and beyond". In M. Timmons (ed.), *Oxford Studies in Normative Ethics* (New York: Oxford University Press) pp. 79–96.
- Shafer-Landau, Russ. (1998). "Moral Judgement and Moral Motivation." *The Philosophical Quarterly* 48 (192): 353–8.

- Shoemaker, David. (2015). *Responsibility from the Margins* (New York: Oxford University Press).
- Smith, Angela. (2005). "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115(2): 236–71.
- Smith, Holly. (2014). "The Subjective Moral Duty to Inform Oneself before Acting." *Ethics* 125: 11–38.
- Smith, Michael. (1994). *The Moral Problem* (Malden: Wiley-Blackwell).
- Sobel, Jordan Howard. (1976). "Utilitarianism and Past and Future Mistakes." *Noûs* 10(2): 195–219.
- Stocker, Michael. (1976). "The Schizophrenia of Modern Ethical Theories." *The Journal of Philosophy* 73(14): 453–66.
- 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.
- Timmerman, Travis and Cohen, Y. (2016). "Moral Obligations: Actualist, Possibilist, or Hybridist?" *Australasian Journal of Philosophy* 94(4): 672–86.
- Vorobej, Mark. (2000). "Prosaic Possibilism." *Philosophical Studies* 97(2): 13–26.
- Wedgwood, Ralph. (2009). "Against Actualism." PEA Soup Blog, URL <http://peasoup.typepad.com/peasoup/2009/09/against-actualism.html>.
- Zimmerman, Michael J. (1996). *The Concept of Moral Obligation* (Cambridge: Cambridge University Press).
- Zimmerman, Michael J. (1997). "A Plea for Accuses." *American Philosophical Quarterly* 34(2): 229–43.
- Zimmerman, Michael J. (2017). "Prospective Possibilism." *Journal of Ethics* 21(2): 117–50.