

Actualism, Possibilism, and the Nature of Consequentialism

Yishai Cohen & Travis Timmerman

Forthcoming in the Oxford Handbook of Consequentialism

Abstract:

The actualism/possibilism debate in ethics is about whether counterfactuals of freedom concerning what an agent would freely do if they were in certain circumstances even partly determines that agent's obligations. This debate arose from an argument against the coherence of utilitarianism in the deontic logic literature. In this chapter, we first trace the historical origins of this debate and then examine actualism, possibilism, and securitism through the lens of consequentialism. After examining their respective benefits and drawbacks, we argue that, contrary to what has been assumed, actualism and securitism both succumb to the so-called nonratifiability problem. In making this argument, we develop this problem in detail and argue that it's a much more serious problem than has been appreciated. We conclude by arguing that an alternative view, hybridism, is independently the most plausible position and best fits with the nature of consequentialism, partly in light of avoiding the nonratifiability problem.

Keywords:

consequentialism, actualism, possibilism, securitism, hybridism, maximalism, omnism, counterfactuals of freedom, action-guiding, nonratifiability problem

1. Introduction

We will use the term *consequentialism* to pick out views in normative ethics which hold that the deontic status of an act-set is solely a function of the net value of the consequences of performing

the act-set in question, relative to the net value of the other act-sets that the agent can perform.^{1,2} Consequentialist views are demarcated from one another by (i) their evaluative rankings of outcomes and, more fundamentally, (ii) their account of the features that determine the evaluative rankings of outcomes. For example, with respect to (i), two consequentialist views may disagree about whether an outcome in which one billion people are prevented from experiencing a headache is better than an outcome in which one innocent person is prevented from dying. With respect to (ii), this disagreement may be the result of each view identifying different intrinsic goods and bads. Perhaps preventing one billion people from experiencing a headache maximizes hedonic utility, while preventing the innocent person from dying maximizes preference satisfaction.³

Category (ii) ranges over a wide variety of issues, though the literature on consequentialism has primarily focused on identifying the correct axiology (e.g. what is intrinsically good and bad), identifying the types of reasons that exist (e.g. moral, prudential), and identifying their respective weights. These distinctions do not exhaust the differences between consequentialist views, however. Any form of consequentialism must also take a stance on the scope of possible acts that are relevant *options* for the agent given certain facts about how an agent would (or might) freely act under certain circumstances. These issues center around the nature of control that an agent must have over their behavior, both in the immediate future and across one's life, and are relatively

¹ Cf. Dreier (2011: 97) and Portmore (2009: 330; 2011: 34-38).

² We will suppose that an agent *S* performs an act-set (singleton or plural) *x* at time *t* iff *x* is a possible state of affairs involving an action (or a number of actions) that belong(s) to *S* and *x* is actualized at *t*.

³ When such cases are discussed in the literature, the focus of the disagreement is about whether there is a lexical priority between values. See Norcross (1998) and Dorsey (2009).

recent in the history of ethics. The question of the scope of an agent's options most notably arises in the actualism/possibilism (A/P) debate, which is the focus of this paper.⁴

The A/P debate may be illustrated with the following case. Suppose that you've just been invited to an ex-partner's wedding, and while you can attend the wedding and be pleasant at little cost to yourself, you're also prone to fits of jealousy. The *best* act-set you can perform involves <accepting the invitation, attending the wedding, and staying sober>, ensuring that everyone has a good time. The *worst* act-set you can perform involves <accepting the invitation, attending the wedding, and getting inebriated>, thereby ruining the wedding for everyone. The act-set that is neither the best nor the worst involves <declining the invitation, and doing something else besides attending the wedding>.⁵ Finally, suppose that if you were to decide to accept the invitation today, then you would, in fact, decide to get inebriated once you're at the wedding.

The question at the root of the A/P debate is whether you are obligated to accept or decline the wedding invitation. This question is particularly important for consequentialists since they must take a stance on whether true counterfactuals of freedom (CFs)—true subjunctive

⁴ These issues also feature centrally in certain debates about deontic logic and the related maximalism/omnism debate, which we'll briefly discuss in section 2.

⁵ This case is drawn from Zimmerman (2006: 153). Since this chapter is focused on the actualism/possibilism debate as it applies to consequentialism, the ranking of outcomes should be understood to be determined by intrinsic value, as opposed to deontic value.

conditionals concerning what an agent *S* would freely do at t_2 if *S* were in circumstances *C* at t_1 —even partly determine an agent’s moral obligation to perform some act-set.^{6,7}

There are four types of views in the debate, viz. possibilism, actualism, securitism, and hybridism, and each of these types admits of multiple distinct precisifications. In response to the central question of the debate, possibilists and some hybridists believe that, necessarily, true CFs do not even partly determine an agent’s obligations, while actualists, securitists, and some hybridists believe that true CFs can at least partly determine an agent’s obligations. So, with respect to the wedding invitation example, actualist forms of consequentialism maintain that you are obligated to decide to decline the invitation at least partly because what would actually happen if you were to decide to decline is better than what would actually happen if you were to decide to accept. Possibilist forms of consequentialism hold that you are obligated to decide to accept because doing so is part of the best act-set you can perform over time. Securitist and hybridist forms of consequentialism are a bit more complex. We will have to consider more detailed versions of the wedding invitation case before we are in a position to explain the implications of these views.

Our aim is to elucidate the differences between actualist, possibilist, securitist, and hybridist forms of consequentialism and, in doing so, explore the benefits and drawbacks of each

⁶ The A/P debate has typically focused on CFs in which ‘*C*’ refers to *S*’s performing an action *x* at t_1 . It also seems to conceive of freedom in terms of having an ability to do otherwise. This is because everyone in the debate implicitly assumes that an agent can do something, such that if they were to do it, then some true CF would be false instead (Cohen & Timmerman 2016: 4).

⁷ The circumstances ‘*C*’ must refer to (temporally intrinsic) maximally specified circumstances at some time in order to accommodate the fact that adding information into the antecedent of a counterfactual can alter its truth-value (Stalnaker, 1968; Lewis, 1973; Jackson, 1985: 178, 186).

from a consequentialist perspective. This chapter is structured as follows. In section 2 we provide some historical context, briefly explaining how a supposed puzzle for utilitarianism within the context of deontic logic gave rise to the actualism/possibilism literature, and then in section 3 we discuss time's crucial relationship to abilities and obligations. In sections 4-8 we review possibilist, actualist, securitist, and hybridist forms of consequentialism in more detail, and argue that hybridism, in one form or other, is both the most plausible view and best captures the nature of consequentialism, partly in light of avoiding the so-called nonratifiability problem as discussed in section 7.

2. Consequentialism and Deontic Logic

The purpose of this section is to trace the historical origins of the A/P debate and its relation to consequentialism. The debate arose partly from Hector-Neri Castañeda's (1968) argument that utilitarianism is formally incoherent. Here's the concise version. Castañeda first assumed that utilitarians accept a principle of deontic logic known as "obligation distributes through conjunction." This principle holds that if an agent *S* is obligated to do both A and B, then *S* is obligated to do A and *S* is obligated to do B (1968: 141). This principle may be represented more formally as follows.

Obligation Distributes Through Conjunction: $O(A \ \& \ B) \rightarrow O(A) \ \& \ O(B)$

This principle, hereafter (ODC), can be illustrated with the wedding invitation case. If you are obligated to <accept the wedding invitation and stay sober at the wedding>, then you are obligated to <accept the wedding invitation> and you are obligated to <stay sober at the wedding>.

Second, Castañeda (1968: 142) argued that (ODC) is inconsistent with the following principle, which he took to be a basic commitment of utilitarianism.

(U): *S* is morally obligated to do *x* in circumstances *C* iff *S*'s doing *x* in *C* will bring about a greater balance of good over bad than her performing any other alternative action open to her in *C*.

To see why (ODC) and (U) are supposedly inconsistent, suppose that an agent *S*'s performing the conjunctive act-set <A & B> brings about a greater balance of good over bad than any alternative act-set (singleton or plural) that *S* can perform. It is supposed to follow from (U) that *S* is obligated to perform <A & B>. Now, given (ODC), it follows that *S* is obligated to perform <A> and that *S* is obligated to perform . But, given (U), performing <A> would result in more net good than performing any alternative, including , *and* performing would *also* result in more net good than performing any alternative, including <A>. But performing <A> cannot result in both more and less net good than performing .⁸

Dag Prawitz (1970) and Fred Westphal (1972) responded to Castañeda by modifying (U) in such a way that the actions are indexed to the time they would need to be performed in order to bring about the uniquely optimific outcome. On their suggested revision, if performing the joint act <A & B> from t_1 - t_2 would result in the greatest net amount of good, then *S* is obligated to

⁸ Lars Bergström (1968: 43) responded to Castañeda's argument against utilitarianism by pointing out that the contradiction arises only if it is assumed that <A> and are *alternatives* in the relevant sense, but <A> and are not, in fact, alternatives since the agent can consistently perform both of them.

perform $\langle A \rangle$ at t_1 and to perform $\langle B \rangle$ at t_2 . This supposedly avoids the contradiction because $\langle A \rangle$ supposedly is *the* act-set at t_1 that would produce the greatest net amount of good in comparison to any other act-set performable at t_1 , whereas $\langle B \rangle$ supposedly is *the* act-set at t_2 that would produce the greatest net amount of good in comparison to any other act-set that is performable at t_2 . Finally, from among the performable act-sets that might occur from t_1 - t_2 , $\langle A \ \& \ B \rangle$ is *the* act-set that would produce the greatest net amount of good in comparison to any other act-set that is performable from t_1 - t_2 . Thus, each act-set is the uniquely optimific one at the time(s) it is performed.⁹

Harold Zellner (1973) demonstrated that indexing actions to times will not solve the supposed problem Castañeda identified for utilitarianism. While performing $\langle A \ \& \ B \rangle$ from t_1 - t_2 may be uniquely optimific, it does not follow that performing either of these individual acts at their respective times would be uniquely optimific. For instance, performing $\langle A \rangle$ at t_1 may not be uniquely optimific if the agent would *not* perform $\langle B \rangle$ at t_2 if she were to perform $\langle A \rangle$ at t_1 . This point may be illustrated with the following case.

Covering Class: The best act-set Bill can perform is $\langle A \ \& \ B \rangle$ from t_1 - t_2 , where $\langle A \rangle =$ agree to teach Ted's class next week, and $\langle B \rangle =$ teach Ted's class next week. Bill can also $\langle C \rangle$ at t_1 , where $\langle C \rangle =$ suggest that George, an inferior instructor, cover Ted's class instead. Finally, suppose that if Bill were to $\langle A \rangle$ agree to teach Ted's class next week, he would $\langle \sim B \rangle$ not teach class.

⁹ Notably, Prawitz (1968; 1970) and Westphal (1972) each argued that an act is permissible if and only if it is part of an act-set that, if performed, would bring about the greatest net good of any of the act-sets available to the agent. Prawitz and Westphal were essentially giving what may be considered the earliest defenses of possibilism, though they did not refer to this view as such.

Again, it would be best if Bill <agrees to teach and teaches>, second best if he <suggests that George cover the class>, and worst if he <agrees to teach and skips teaching>. Thus, the value of the act-sets may be ranked from best to worst as follows.

(1) <A & B>

(2) <C >

(3) <A & ~B>

Zellner points out that, since Bill would <~B> if he were to <A>, the value of performing <A> is not uniquely optimific at t_1 even though the value of performing <A & B> from t_1 - t_2 is uniquely optimific.¹⁰ Thus, in such cases, (U) combined with (ODC) still generates contradictions, even if each act-set (singleton or plural) is indexed to their respective times.

Zellner's own proposed solution was to reject (U) because it is inconsistent with a supposedly basic principle of inference which holds that if an agent is obligated to perform <A> and her performing <A> entails her performing , then she is obligated to perform (1972: 125). This rule, sometimes referred to as *Normative Inheritance* (NI), may be represented more formally as follows (Feldman 1986: 41).

Normative Inheritance: If $\vdash A \rightarrow B$ then $\vdash O(A) \rightarrow O(B)$

¹⁰ An action that is uniquely optimific at a time is one that is optimific relative to all other performable actions at that time.

Regardless of whether Zellner's own response to Castañeda succeeds, cases such as *Covering Class* highlight the important question within the A/P debate as to whether true CFs even partly determine an agent's obligations.

Since a conjunction entails each of its conjuncts, it follows that (NI) entails (ODC). Moreover, a commitment to (NI), and thus to (ODC), suggests that, in addition to an entailment relation, there is also a dependence relation between obligatory act-sets. We can distinguish between a *non-dependent* obligation which we have not in virtue of having some other obligation and a *dependent* obligation which we have in virtue of having some other obligation. All non-dependent obligations are ultimately possessed in virtue of possessing some non-dependent obligation (Timmerman & Cohen, 2016: 679; cf. Portmore, 2011: 179).

Many such theories that are committed to (ODC) or (NI) are forms of maximalism. They maintain that, necessarily, the object of a non-dependent obligation for an agent *S* is a maximal act-set, or something close enough.¹¹ A maximal act-set is roughly one that, at some time *t*, *S* can perform across their entire life up to the last performable act. More precisely, a maximal act-set is an act-set that, at some time *t*, an agent *S* can perform over time, and it is not contained in some other act-set that, at *t*, *S* can perform over time. Act-set *x* is contained in act-set *y* iff (i) *x* and *y* belong to the same agent, (ii) the period of time at which *x* is performed is a proper or improper part of the period at which *y* is performed, and (iii) it is logically necessary that if *y* is performed

¹¹ We agree with Portmore and Brown that, more exactly, the object of a non-dependent obligation is a maximally normatively specific option (or act-set). This is an act-set that is entailed only by normatively equivalent options/act-sets (Portmore, 2017: 428; Brown, 2018: 13). All of the arguments put forward in this paper regarding the status of maximal act-sets similarly apply, *mutatis mutandis*, to the status of maximally normatively specific options/act-sets.

then x is performed (Prawitz, 1968: 80; Sobel, 1976: 199; Portmore, 2011: 177; Brown, 2018: 752). The final crucial component to maximalism is that every non-maximal act-set that is contained in a non-dependent, obligatory maximal act-set is itself the object of a dependent obligation.

In contrast to maximalism, omnism implies that the object of a non-dependent obligation may be a maximal or non-maximal act-set because the deontic status of every act-set x , maximal and non-maximal, is to be evaluated only in terms of its own value rather than in terms of its relation to the value of an act-set y that entails x or which contains x (Portmore, 2017: 429, 431). With the exception of the version of actualism to be discussed in section 5, all of the theories to be discussed presuppose maximalism, or something similar to maximalism (see footnote 11).

In the next section, we provide an elaborated version of the wedding invitation case that will guide us through the A/P terrain, and we will highlight a number of background assumptions that we explicitly adopt in our approach to the debate.

3. Time's Relationship to Abilities and Obligations

Before we can explore the views of the A/P debate in detail, we must first discuss time's role with respect to an agent's abilities and obligations. We'll do so, in part, through further discussion of the aforementioned principles of deontic logic. This will provide the background information necessary to understand the subsequent detailed discussion of each of the views in the A/P debate.

Let's start by considering an elaborated version of the wedding invitation case, which we'll refer to as *Wedding Invitation 1 (W11)*. Suppose that Alice and her ex-partner promised each other in the past that they would attend each other's weddings, and that Alice has been invited to her ex-

partner's wedding. At t_1 Alice is <A>: deliberating about what to do. The following actions are ones that, at t_1 , she can perform in the future:

Possible actions at t_2 :

B = Decide to: accept the wedding invitation, attend the wedding and not drink alcoholic beverages from t_2 - t_4

C = Decide to: decline the wedding invitation at t_2 , and then go home and do research for a paper from t_3 - t_4

Possible actions at t_3 :

D = Decide to: attend the wedding and not drink alcoholic beverages from t_3 - t_4

E = Decide to: go home and play videogames from t_3 - t_4

F = Decide to: go home and do research for a paper from t_3 - t_4

G = Decide to: go home and play videogames from t_3 - t_4

Possible actions at t_4 :

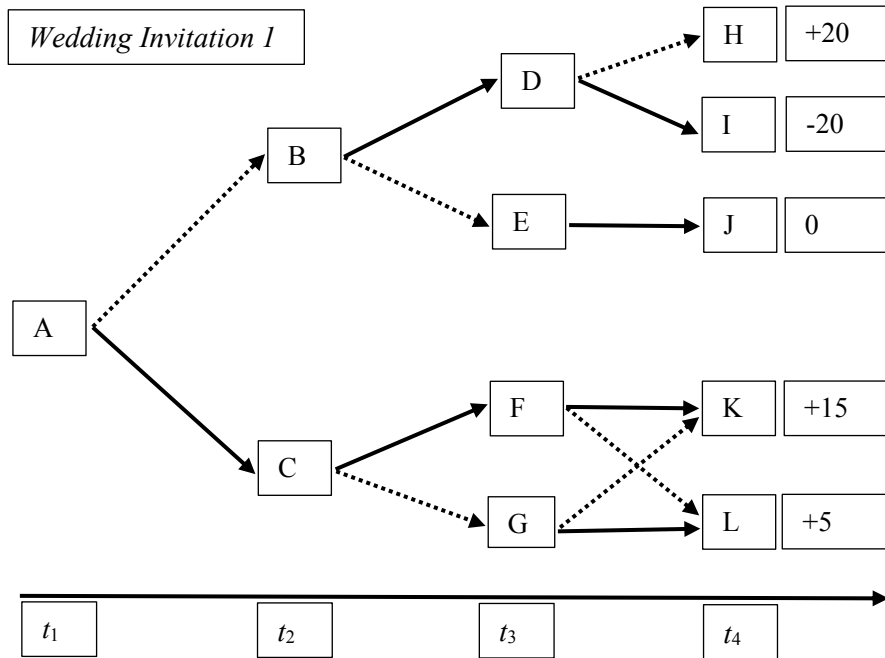
H = Decide to: not drink alcoholic beverages at t_4

I = Decide to: drink an alcoholic beverage at t_4

J = Decide to: play videogames at t_4

K = Decide to: do research for a paper at t_4

L = Decide to: play videogames at t_4



All arrows point to what Alice *can* do at the respective times, but only the solid arrows indicate what Alice *would* subsequently do if she were in certain prior circumstances. Notice, then, that, no matter what Alice decides at t_2 to do from t_2 - t_4 among the decisions that, at t_1 , she can make at t_2 , it's not the case that Alice would perform any of the actions (more specifically, decisions) that are preceded at any point in time by a dotted line. The positive and negative numbers represent the total net hedonic utility (i.e. the total number of hedons and dolors) of the consequences that would be brought about by each of the above act-sets. To keep things simple, let's suppose that hedonistic act-utilitarianism is true, though, to be clear, all arguments put forward in this chapter can accommodate any version of consequentialism, including those that posit agent-relative constraints and/or pluralism about intrinsic value. Let's also suppose for simplicity's sake that, for whatever reason, the final act that Alice can perform over time extends no later than t_4 .

At t_1 , $\langle B, D, H \rangle$ is the best maximal act-set that Alice can perform over time, $\langle B, D, I \rangle$ is the *worst* maximal act-set that Alice can perform over time, and $\langle C, F, K \rangle$ is the *second-best* maximal act-set she can perform over time. As a matter of fact, Alice performs the second-best

maximal act-set that, at t_1 , she can perform from t_2 - t_4 . Moreover, if Alice were to $\langle B \rangle$ at t_2 , then she would, in fact, $\langle D \rangle$ at t_3 and $\langle I \rangle$ at t_4 .

With this case in mind, let's consider what Alice's abilities are at any given time. It is natural, we think, to understand an agent's *present* ability to do something as an ability to bring about a *future* event that might occur rather than as an ability to bring about a present event that is already happening, though many have suggested that this distinction doesn't have any significant ramifications for the A/P debate.¹² We will adopt an exclusively future-oriented approach, according to which an agent has a present ability to settle which of the alternative *future* options will be realized.

4. Possibilism

Now that we have reviewed the relevant background information, we are in a position to dive deeper into the A/P debate by considering each of the views in more detail. In this section, we will examine possibilism and explain why it is viewed as an attractive position, while also reviewing the most compelling objection against it. Here is a more formal definition of possibilism.

Possibilism: At t an agent S is obligated to φ at t' iff S can, at t , φ at t' and φ -ing at t' is part of the best act-set that, at t , S can perform from t' until the last time she can perform an act.

According to possibilism, an agent's sole *non-dependent obligation* is to perform an optimistic *maximal act-set* (i.e. the best act-set agents can perform over the course of their entire lives), and they will have *dependent obligations* to perform the non-maximal act-sets that are contained in

¹² Cf. Carlson (1995: 77), Goldman (1976: 453), Bykvist (2002: 47), and Portmore (2011: 166).

this non-dependent, obligatory maximal act-set. In *WII*, for instance, possibilists hold that Alice has a dependent t_1 -obligation to $\langle B \rangle$ at t_2 in virtue of having a non-dependent t_1 -obligation to $\langle B, D, H \rangle$ from t_2 - t_4 .¹³

Possibilism has a number of defenders for a reason. In addition to generating the intuitively correct moral verdicts in a wide range of cases, it also comports well with consequentialist commitments. Specifically, possibilism captures the most straightforward understanding of the idea that agents are obligated to bring about the best outcome they can. It's for this reason that consequentialists¹⁴ seem to frequently implicitly assume possibilism in their work. A second set of advantages is that possibilism preserves a number of plausible axioms in deontic logic, most notably (ODC) and (NI).¹⁵ A third set of advantages is that possibilism avoids the primary objections to actualism to be discussed in the next section.¹⁶ While possibilism no doubt has a lot going for it, it also is subject to a variety of objections, one of which is particularly notable.

¹³ The ability to perform a (non-singleton) act-set over time may be understood in terms of having an ability to perform the first act in the act-set, and then, once the first act is performed, they will have the ability to perform the second act in the act-set, and once the second act is performed, they will have the ability to perform the third act, etc., culminating in the agent's ability to perform the final act in that act-set. See e.g. Goldman (1978: 193) and Portmore (2011: 165-166, 170).

¹⁴ This notably includes consequentialist effective altruists. See Timmerman (in press).

¹⁵ Goldman, 1978: 80; Feldman, 1986: 41-44; Zimmerman, 1990: 58-60; Zimmerman, 2006: 154-155; Vessel, 2009; Kiesewetter, 2018.

¹⁶ There is more to be said in favor of possibilism that is beyond the scope of this paper. See Zimmerman (1996: fn. 72 and fn. 122) and (2017: ch. 3) for a nice review of some of possibilism's additional, lesser appreciated, virtues.

The most influential objection to possibilism is that it can generate an obligation that, if acted on, would result in the worst possible outcome.¹⁷ Possibilism generates this consequence because it implies that true CFs do not even partly determine an agent's obligations. This objection may be formulated more precisely as follows.

Worst Outcome Objection: Possibilism entails that an agent *S* can have an obligation to φ even when φ -ing would result in *S* performing an act-set that is deeply morally wrong (perhaps the worst possible act-set) and that is worse than the act-set that *S* would perform if *S* were to $\sim\varphi$.

To illustrate, possibilism says that in *WII* Alice has a dependent t_1 -obligation to $\langle B, D \rangle$ from t_2 - t_3 in virtue of having a non-dependent t_1 -obligation to $\langle B, D, H \rangle$ from t_2 - t_4 . However, if Alice were to $\langle B, D \rangle$, then she would $\langle I \rangle$ rather than $\langle H \rangle$ at t_4 , resulting in the worst possible outcome. This may not sound so counterintuitive in cases where the worst outcome isn't tragic. Yet, this objection has more intuitive force in high-stakes variants. Suppose that, no matter what Alice decides prior to t_4 to do, if Alice were to drink alcohol at the wedding, then she would drive home drunk, killing three pedestrians in the process. Possibilism still renders the verdict that Alice has a non-dependent t_1 -obligation to $\langle B, D \rangle$ from t_2 - t_3 , and it renders this verdict no matter how terrible the consequences of doing this happen to be. So, while possibilism can account for consequentialists' judgments that agents are obligated to bring about the best outcome that they can over time, they

¹⁷ Goldman, 1976: 469-70; Sobel, 1976: 202-203; Feldman, 1986: 52-57; Almeida, 1992: 461-462; Woodard, 2009: 219-221; Ross, 2013: 81-82; Gustafsson, 2014: 593; Timmerman & Cohen, 2016: 674.

cannot account for consequentialists' judgments that agents are always obligated to act in ways that *would* bring about the best outcome.

Possibilists have tried to sugar the pill of the *Worst Outcome Objection* by emphasizing the distinction between conditional and unconditional obligations. They'll note that the possibilist obligation picks out agents' unconditional obligations and still allows that agents can incur conditional obligations to act in ways that preclude them from bringing about the worst outcome.¹⁸ So, while it's true that possibilism entails that Alice has an unconditional obligation to in virtue of her unconditional obligation to <B, D, H>, it may also be true that she has a conditional obligation to <C> given that Alice would <I> if she were to . More generally, possibilists respond to this objection by claiming that agents have an unconditional obligation to do the best they can but incur conditional obligations to bring about the next best outcome *if* they won't bring about the best outcome. While the appeal to conditional obligations renders possibilism more palatable, the *Worst Outcome Objection* nevertheless reveals an important tension between possibilism and consequentialist commitments.

5. Actualism

In this section, we will examine a standard form of actualism and then review two problems with actualism that concern the relationship between control and the truth-value of certain CFs. Here is a standard definition of actualism.

¹⁸ See Greenspan (1978: 81) and Zimmerman (1986: 70; 2017: 126-128).

Actualism: At t an agent S has an obligation to φ at t^* iff, at t , S can φ at t^* ($t < t^*$), and what would happen if S were to φ at t^* is better than what would happen if S were to perform any other act-set that, at t , S can perform at t^* .¹⁹

This standard definition of actualism implies that Alice has a t_1 -obligation to $\langle C \rangle$ at t_2 , not because it is contained in the best act-set that, at t_1 , Alice can perform over time, but rather because what would happen if Alice were to $\langle C \rangle$ at t_2 is better than any other act that, at t_1 , Alice can perform at t_2 . Actualism thus avoids the worst outcome objection precisely by maintaining that true CFs at least partly determine an agent's obligations.

One problem with this definition is that it does not require φ -ing at t^* to be a *fully specified* act-set, i.e. an act-set that is not contained in any other act-set that, at t , S can perform at t^* .²⁰ To see why ' φ ' needs to be fully specified, consider another wedding scenario, *WI2*, that has the following deviations from *WI1*.

Alice's neighbors always have loud and disruptive parties, and, rather than $\langle K \rangle$ and $\langle L \rangle$ being options for Alice at t_3 , Alice can, at t_3 , perform any of the following acts at t_4 :

K1 = \langle Decide to: do research and kill the neighbors at t_4 \rangle

K2 = \langle Decide to: do research and *not* kill anyone at t_4 \rangle

¹⁹ This formulation of actualism avoids what is often referred to as the lumping problem (Wedgwood, 2009; Cariani, 2016). A different formulation of actualism may be found in Jackson & Pargetter (1986). For a detailed overview of their version, and how it is different from alternative versions, see Timmerman & Cohen (2019).

²⁰ Note that all maximal act-sets are fully specified act-sets, but not all fully specified act-sets are maximal act-sets. For example, suppose that an agent has, at t_1 , an ability to perform a fully specified act-set $\langle x \rangle$ at t_2 . Hence, there is no act-set that is not identical to $\langle x \rangle$ that, at t_1 , the agent can perform at t_2 that contains $\langle x \rangle$. However, $\langle x \rangle$ itself is not maximal if it is contained in the following maximal act-set that the agent can, at t_1 , perform *over time* from t_2 - t_3 : $\langle x, y \rangle$.

L1 = <Decide to: play videogames and kill anyone at t_4 >

L2 = <Decide to: play videogames and *not* kill anyone at t_4 >

Suppose that, as a matter of stipulation, Alice would <L2> at t_4 if she were to <F> at t_3 , and that if doing research were part of the content of Alice's decision at t_4 then killing the neighbors would also be part of the content of Alice's decision at t_4 (the closest world in which Alice does research is one in which she kills the neighbors). Nevertheless, in light of the fact that, at t_3 , Alice can <K2> at t_4 , we may conclude that Alice is still obligated to form a decision whose content includes doing research since <K2> is the best fully-specified act-set at t_4 that, at t_3 , Alice can bring about. The lesson to be gleaned from this, according to Goldman (1978:186-190), is that no agent has an obligation even partly in virtue of the truth-value of what we will call a *synchronic* CF—i.e. a counterfactual in which the antecedent and the consequent are both indexed to the same time. So, the following true synchronic CF does not even partly determine Alice's obligations at t_3 : 'If Alice were to <do research> at t_4 , then Alice would also <kill the neighbors> at t_4 '.²¹ This is why the object of an agent's obligation in any A/P scenario must be a fully specified act-set. So, we will suppose that each of the act-sets mentioned in *WII* is fully specified and that the content of each decision implicitly excludes the performance of any other number of normatively significant acts like killing the neighbors or donating to charity. Goldman's remarks thus reveal that the kind of control an agent must have in order to have a moral obligation involves a control over the truth-value of the relevant synchronic counterfactuals.

²¹ This counterfactual is true since, as a matter of stipulation, both its antecedent and its consequent are true (Pollock, 1976: 42–43; Stalnaker, 1968; Lewis, 1986: 132).

Another kind of counterfactual that does not even partly determine an agent's present obligations concerns what an agent would subsequently freely do if they were in the very circumstances in which they presently find themselves. For instance, when considering what, at t_1 , Alice ought to do at t_2 , all parties in the debate assume that what Alice would freely do at t_2 if she were to $\langle A \rangle$ at t_1 does not even partly determine Alice's t_1 -obligations. We will call this kind of counterfactual an *early* counterfactual because the antecedent involves a time t at which an agent has an obligation to subsequently do something at t^* ($t < t^*$). No party in the A/P debate assumes that early counterfactuals even partly determine an agent's obligations. To illustrate this point further, consider a scenario, *WI3*, that is similar to *WI1*, except that Alice performs $\langle B, D \rangle$ from t_2 - t_3 . Actualists and possibilists alike agree that, at t_3 , Alice ought to $\langle H \rangle$ at t_4 , and so the following true CF does not even partly determine Alice's t_3 -obligation to do something at t_4 : 'If Alice were to $\langle D \rangle$ at t_3 , then Alice would $\langle I \rangle$ at t_4 '. After all, the truth of this early counterfactual is incompatible with Alice's performing $\langle H \rangle$ at t_4 , and at t_3 Alice can do something at t_4 , viz. $\langle H \rangle$, such that if she were to do it, then this early counterfactual would be false (Cohen & Timmerman 2016). Having control over the truth-value of an early counterfactual highlights the importance of an agent's control over their immediately available action(s). A more plausible position in the A/P debate must take these insights into account, and securitist views of different stripes do exactly that.

6. Securitism

Securitist views may be thought of as versions of actualism that took these insights into account and were modified to focus on the best outcome that would occur from among all of the fully specified act-sets that an agent can immediately perform. In this section, we will briefly review the

most important kinds of securitism in the literature before reviewing two of the most influential objections that apply to each view. Perhaps the most popular kinds of securitism are maximalist versions, which hold that an agent's options at t are the jointly exhaustive and mutually exclusive, fully specified, maximal act-sets that are securable for an agent at t .²² Here is a formal definition of maximalist securitism.

Securitism: At t an agent S has a non-dependent obligation to φ at t^* iff, φ is a fully specified, maximal act-set that, at t , is securable for S at t^* ($t < t^*$), and what would happen if S were to φ at t^* is better than what would happen if S were to perform any other fully specified, maximal act-set that, at t , is securable for S at t^* .

A maximal act-set L ²³ is securable for S at time t if, at t , S can immediately perform the first moment of x , x is in L , and if x were to occur, then L would occur (Sobel 1976: 199). Not all forms of securitism are maximalist, however. Goldman's version of securitism is neither maximalist nor omnist. Instead, Goldman (1978: 194-195) maintains that, at t , an agent can (in the relevant sense of "can") perform an act-set over time only if, at t , an agent can form an intention to perform this act-set, the act-set is securable in virtue of the decision being causally efficacious if performed in the sense that the content of the decision would be actualized if the decision were to occur, and

²² See Goldman (1978: 202) and Portmore (2011: 179). Gustafsson's (2014) solution to Castañeda's (1968) formal critique of consequentialism appeals to jointly exhaustive and mutually exclusive act-sets. Gustafsson considers his position to be distinct from maximalism, but Brown (2018: 66-67) persuasively argues that Gustafsson's position is consistent with maximalism.

²³ Sobel (1976) refers to a maximal act-set as a "life".

this content corresponds to the relevant securable act-set.²⁴ We stipulated in *WII* that the content of the decisions involve all of the relevant activity up to t_4 in order to make Goldman's verdict about *WII* align with the verdict of maximalist versions of securitism, such as Portmore's version. Both hold that Alice has a non-dependent t_1 -obligation to $\langle C, F, K \rangle$ from t_2 - t_4 , and that this act-set is securable for Alice at t_1 in virtue of Alice having, at t_1 , the ability to perform the fully specified singleton act-set $\langle C \rangle$ at t_2 .

We now turn to the most important objections in the literature that apply to both actualism and securitism. The first influential objection is that both views let agents off the hook too easily by allowing them to avoid incurring moral obligations in virtue of their rotten moral dispositions.²⁵ In *WII*, for instance, both views allow Alice to avoid incurring an obligation to attend the wedding simply because she is disposed to act wrongly if she were to freely decide to accept the invitation. With respect to securitist views, it matters that it is not securable at t_1 for Alice to both attend the wedding and refrain from drinking (Timmerman 2015; Vessel 2016). Nevertheless, at t_1 , attending the wedding is easily securable for her, and if Alice were to attend, then not drinking would be easily securable for her at t_3 .²⁶ But, it seems, agents shouldn't get out of having to do good things just because they're disposed to do bad things. This objection may be stated more precisely as follows.

²⁴ Cf. Bykvist (2002: 50-51).

²⁵ However, see Timmerman & Swenson (2019) for an argument that possibilism is subject to an analogue problem.

²⁶ We take it that some actions are more difficult to perform than others. Fill out this account whichever way you like, and suppose that each securable act is easy to perform in this sense.

The Not Demanding Enough Objection: Actualism and securitism permit an agent *S* to avoid incurring any moral obligation to ϕ , which *S* can easily fulfill, simply in virtue of *S*'s rotten moral character.²⁷

The second influential objection is that both views prescribe the agent to perform harmful actions even when the agent can refrain from performing any such harmful actions. To take an extreme example, in certain cases they could require someone to murder five innocent people if they would otherwise murder six innocent people even though (everyone agrees) the agent can refrain from murdering anyone. In the less extreme example of *WII*, both views entail that, at t_1 , Alice has a dependent obligation to $\langle C \rangle$ at t_2 and that, as a result, actualism and securitism prescribe the bad behavior of breaking a promise to attend someone's wedding. That's not quite as bad as murdering five innocent people, but it still seems objectionable given that Alice can attend the wedding and be pleasant. This objection may be formulated more precisely as follows.

The Bad Behavior Objection: Actualism and securitism prescribe bad behavior, and acting on such prescriptions presumably renders²⁸ an agent *S* immune from moral criticism, even when *S* can easily refrain from such behavior.²⁹

²⁷ Jackson & Pargetter, 1986: 240; Zimmerman, 2006: 156; Portmore, 2011: 207; Baker, 2012: 642-3); Timmerman, 2015: 1,512-1,513.

²⁸ This is not strictly entailed by actualism, but it is entailed by actualism coupled with widely accepted axioms about moral blameworthiness.

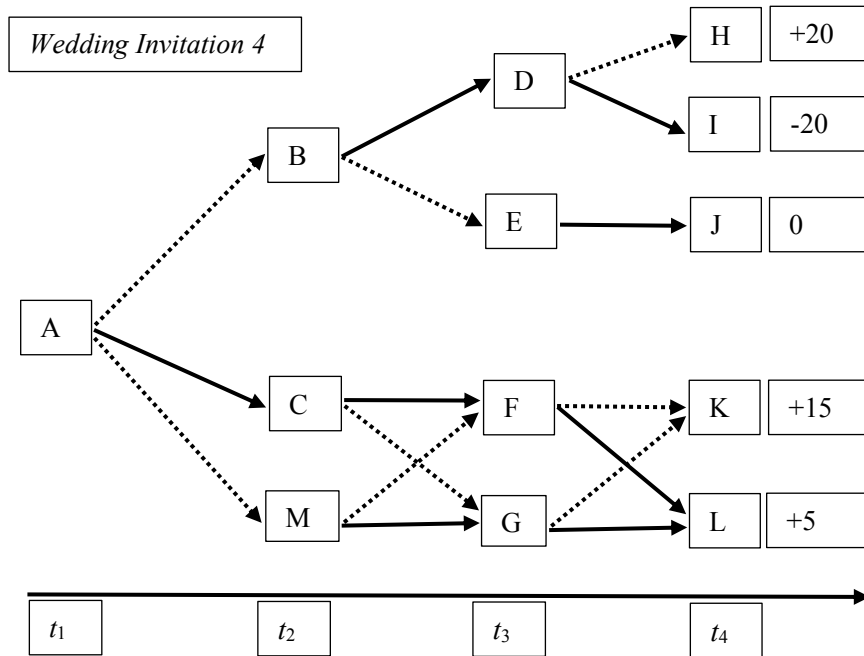
²⁹ Wedgwood, 2009; Ross, 2013; Timmerman & Cohen, 2016; Zimmerman, 2017: 121.

In the next section we analyze an underexplored objection, the nonratifiability problem, and argue that, contrary to what has been assumed, all versions of actualism and securitism are subject to this problem. This will also help motivate hybridism, which we discuss in section 8.

7. The Nonratifiability Problem

Securitists such as Goldman (1978: 202) and Portmore (2011: 179, 181-182), maintain that, in *WII*, Alice has at least three dependent t_1 -obligations to $\langle C \rangle$ at t_2 , $\langle F \rangle$ at t_3 , and $\langle K \rangle$ at t_4 , respectively, and that Alice has these obligations in virtue of having a non-dependent t_1 -obligation to $\langle C, F, K \rangle$ from t_2 - t_4 . Moreover, both agree that Alice has a (non-dependent) t_3 -obligation to $\langle K \rangle$ at t_4 , which is exactly what we would expect given that Alice violates no obligation (according to both actualism and securitism) from t_2 - t_3 , and given that Alice has a (dependent) t_1 -obligation to $\langle K \rangle$ at t_4 . However, here's a problem for such views. Fulfilling one's obligations during some period of time does *not* guarantee a consistency among an agent's various dependent singleton obligations across time. To see the problem, consider a revised version of the wedding case, *WI4*, in which Alice performs $\langle L \rangle$ rather than $\langle K \rangle$ at t_4 , and so the following CF that is false in *WII* is true in *WI4*: 'If Alice were to $\langle F \rangle$ at t_3 , then she would $\langle L \rangle$ at t_4 '. Let's also suppose that Alice has the following additional option in our new case, *WI4*:

M = Decide to: decline the wedding invitation at t_2 , and then go home and play videogames from t_3 - t_4



WI4 reveals that both actualism and securitism allow for a conflict among an agent's obligations across time, even when an agent never violates an obligation during that interval of time. Both views hold that Alice has a dependent t_1 -obligation to $\langle L \rangle$ at t_4 in virtue of having a non-dependent t_1 -obligation to perform a maximal act-set that contains $\langle L \rangle$. Alice has this non-dependent t_1 -obligation because such an act-set is a better securable one for Alice at t_1 than any other securable act-set for Alice at t_1 that does not contain $\langle L \rangle$, viz. $\langle B, D, I \rangle$ ($+5 > -20$). Nevertheless, regardless of whether Alice performs $\langle F \rangle$ or $\langle G \rangle$ at t_3 , Alice has a t_3 -obligation to $\langle K \rangle$ rather than $\langle L \rangle$ at t_4 ($+15 > +5$), even though $\langle L \rangle$, and not $\langle K \rangle$, is one of Alice's dependent t_1 -obligations. Ross (2012: 87-89) calls this phenomenon the nonratifiability problem: actualism and securitism make nonratifiable prescriptions, i.e. prescriptions that will inevitably be reversed prior to the time of the act, no matter what the agent does prior to that act.

There is a specific sense in which actualism and securitism appear *not* to be action-guiding in light of these conflicting obligations across time; it is difficult to see in what sense Alice has a t_1 -obligation to $\langle L \rangle$ at t_4 since, no matter what she does prior to t_4 —and, we would like to

emphasize, even if she never violates an obligation prior to t_4 —Alice lacks a t_3 -obligation to $\langle L \rangle$ at t_4 . Possibilism, by contrast, faces no such problem: just as Alice has a dependent t_1 -obligation to $\langle H \rangle$ at t_4 (in virtue of having a non-dependent t_1 -obligation to $\langle B, D, H \rangle$ from t_2 - t_4), Alice would similarly have a t_3 -obligation to $\langle H \rangle$ at t_4 if Alice were to $\langle B, D \rangle$ from t_2 - t_3 .

We think that the nonratifiability problem raises additional unexplored difficulties for actualism and securitism. But before we explore such difficulties, we will first turn to Ross's theory and explain how it aims to avoid the nonratifiability problem. We will argue that, contrary to what Ross claims, his theory is not immune to the nonratifiability problem. Here is Ross's (2012: 91) version of securitism.

Momentwise Wide-scope Securitism (MWSS): For any x and t , at t , x ought to be such that, for all t' from t forward, x satisfies the following conditional: For all Φ , if whether x Φ s does not causally depend on the intentions x has after t' , and if every maximally preferable option that is directly securable for x at t' involves Φ -ing, then x Φ s.

Like *WI4*, Ross considers a case that highlights the nonratifiability problem. Here is a summary of the case (Ross 2012: 87-88):

Satan's School for Girls On June 6 2011 Sally is about to be kidnapped by Satanists and incorporated into Satan's School for Girls. On June 6 2016 Sally can kill her firstborn child (for the prince of darkness) either by cutting off her child's head or by bludgeoning her child to death with an axe, and Sally can also refrain from killing her child. However, while both cutting off her child's head in 2016 and killing her child with an axe in 2016 are

securable for Sally in 2011, refraining from killing her child in 2016 is not securable for Sally in 2011. In other words, even if Sally were to decide in 2011 to refrain from killing her child in 2016, she would still kill her child in 2016.

The best maximal act-set that is securable for Sally on June 6, 2011 involves cutting off her child's head on June 6 2016 (this form of death is less painful). However, at the relevant time on June 6, 2016, the best maximal act-set that is securable for Sally involves refraining from killing her child in 2016. Like *WI4*, we may wonder in what sense Sally has a 2011-obligation to kill her child in 2016 since, even if Sally fulfills all of her obligations up to June 6 2016, Sally will not be obligated to follow through on her 2011-obligation to cut off her child's head. Now, Ross claims that, unlike securitism, *MWSS* avoids the nonratifiability problem:

MWSS implies that, on June 6, 2011, Sally is under an obligation to refrain from killing her firstborn child five years later. For *MWSS* implies that, on June 6, 2011, Sally ought to be such that, *on June 6, 2016*, she satisfies the conditional (if whether Sally refrains from killing her firstborn child on June 6, 2016 does not depend on her intentions *after that time* [i.e. after a point in time on June 6, 2016], and if every maximally preferable option that is directly securable for Sally *on June 6, 2016* involves refraining from killing her firstborn child, then she refrains from killing her firstborn child). And, given the description of the case, come what may, Sally will satisfy the antecedent of this conditional. Thus, the only way Sally can satisfy this conditional is by refraining from killing her firstborn child on June 6, 2016. *MWSS* therefore implies that, on June 6, 2011, Sally is under an obligation to refrain from killing her firstborn child five years later (Ross 2012: 92; italics added).

The first thing to notice about *MWSS* is that the times “for all t ’ from t forward” refer to the times when an agent has an obligation and that the obligatory act, Φ -ing, is implicitly also indexed to a time. It’s important to be clear about this distinction, then, in order to accommodate the fact that we may have the same kind of obligation at *different* times to do something at some future time. For example, possibilism, actualism, and securitism agree that in *WII* Alice has both a t_2 -obligation and a t_3 -obligation to $\langle K \rangle$ at t_4 .

Now, Ross’s claim that *MWSS* avoids the nonratifiability problem relied on a claim about Sally’s 2016-obligation (notice the parts that we italicized), notwithstanding the fact that Ross says “on June 6, 2011, Sally ought to be such that, on June 6, 2016” which is not captured in the definition of *MWSS*, nor does it make any difference to what is being claimed since it is true at *all* times that Sally has a 2016-obligation to refrain from killing her firstborn child in 2016.³⁰

When we turn to ask what, according to *MWSS*, Alice’s 2011-obligation is, *MWSS* seems to imply that, on June 6, 2011, Sally is obligated to cut off her child’s head on June 6, 2016. This is because, *at all times*, Sally ought to be such that on June 6, 2011 she satisfies the following conditional: (if whether Sally refrains from killing her firstborn child on June 6, 2016 does not depend on her intentions after June 6, 2011, and if every maximally preferable option that is directly securable for Sally on June 6, 2011 involves refraining from killing her firstborn child, then she refrains from killing her firstborn child). On June 6, 2011 Sally satisfies this conditional by failing to satisfy the second conjunct of the antecedent: no maximally preferable option that is directly securable for Sally on June 6, 2011 involves refraining from killing her firstborn child.

³⁰ A bit more precisely, Sally has an obligation at t (which is in 2016) to decide at some time t^* after t (which is also in 2016) to refrain from killing her firstborn child (in 2016).

Since Sally need not satisfy the conditional by satisfying the consequent, *MWSS* does not imply that, on June 6, 2011, Sally is obligated to refrain from killing her firstborn child on June 6, 2016. Instead, *MWSS* seems to imply that, on June 6, 2011, Sally is obligated to kill her firstborn child on June 6, 2016. This is because Sally ought to be such that on June 6, 2011 she satisfies the following conditional: (if whether Sally kills her firstborn child on June 6, 2016 does not depend on her intentions after June 6, 2011, and if every maximally preferable option that is directly securable for Sally on June 6, 2011 involves killing her firstborn child, then she kills her firstborn child). Since the antecedent of this conditional is true, the only way for Sally to satisfy this conditional is by satisfying the consequent. Hence, *MWSS* seems to imply that, on June 6, 2011, Sally *is* obligated to kill her child on June 6 2016. So, like securitism, *MWSS* similarly succumbs to the nonratifiability problem. In order to cement this conclusion, let's consider what *MWSS* says about *WI4*.

MWSS says that Alice ought to be such that, at t_2 , she satisfies the following conditional: (if whether Alice <L-s> at t_4 does not depend on her intentions after t_2 , and if every maximally preferable option that is directly securable for Alice at t_2 involves <L-ing> at t_4 , then she <L-s> at t_4). Since Alice satisfies the antecedent of this conditional, the only way to satisfy the conditional is by satisfying its consequent. So, *MWSS* implies that Alice has a t_2 -obligation to <L> at t_4 . However, *MWSS* also implies that Alice has a t_3 -obligation to <K> rather than <L> at t_4 because Alice satisfies the antecedent of the following conditional, (if whether Alice <K-s> at t_4 does not depend on her intentions after t_3 , and if every maximally preferable option that is directly securable for Alice at t_3 involves <K-ing> at t_4 , then she <K-s> at t_4), and so the only way to satisfy the conditional is by satisfying the consequent. So, it seems that, like actualism and other forms of securitism, Ross's version of securitism is, in fact, subject to the nonratifiability problem.

The nonratifiability problem is, we believe, underexplored in the literature. This may be because actualists and securitists regard the problem as a quirky implication of their view, a small bullet to bite at best. However, we believe the nonratifiability problem points to a much deeper issue. We will argue that any view subject to the nonratifiability problem either violates ‘ought’ implies ‘can’ (OIC) or is committed to an implausible position on the relationship between dependent and non-dependent obligations. As such, any view subject to the nonratifiability problem should be rejected.

Consider a case, *WI5*, that is just like *WI4*, except that we add the following stipulations: Alice is a fully rational agent from t_1 - t_2 ,³¹ and she knows from t_1 - t_2 that, no matter what happens up until t_3 , she won’t have a t_3 -obligation to $\langle L \rangle$ at t_4 . Given these stipulations, it follows that Alice cannot, at t_1 , $\langle M \rangle$ at t_2 if we embrace the following principle:

(The Rationality-Ability Principle) A fully rational agent S cannot, at t , decide at t^* ($t < t^*$) to perform an action x at t^{**} ($t^* < t^{**}$) if S knows from $t - t^*$ that, no matter what happens up until t^{**} , S will have, after t^* and prior to t^{**} , decisive reason to refrain from x -ing at t^{**} .

Given the stipulations of *WI5* and the truth of the Rationality-Ability Principle, it follows in *WI5* that, at t_1 , Alice cannot $\langle M \rangle$ at t_2 because part of the content of $\langle M \rangle$ includes playing videogames at t_4 , and Alice knows that, no matter what happens up until t_3 , Alice lacks a t_3 -obligation to play

³¹ One might object that Alice must not be a fully rational agent from t_1 - t_2 in *WI5* because she performs $\langle C \rangle$ at t_2 , and, for some reason, this is not what a fully rational agent would do in these circumstances. Even if this is right, we can add the further stipulation that *WI5* remains silent on what Alice in fact does at t_2 .

(and decide to play) videogames at t_4 . Despite Alice's inability at t_1 to $\langle M \rangle$ at t_2 , it seems at first glance that securitism is committed to the position that Alice has a t_1 -obligation to $\langle M \rangle$ at t_2 which would violate OIC. But suppose the securitist says in response that, since OIC is obviously true, it must not be the case that Alice has such a t_1 -obligation after all. Then what, we may ask, is Alice obligated at t_1 to do at t_2 according to securitism?

Can securitists claim that Alice has a t_1 -obligation to $\langle C \rangle$ at t_2 ? If Alice has a dependent t_1 -obligation to $\langle C \rangle$ at t_2 , this must be in virtue of Alice's non-dependent t_1 -obligation to $\langle C, F, L \rangle$ from t_2 - t_4 , even though the content of $\langle C \rangle$ includes doing research for a paper and *not* playing videogames. However, despite doing the right thing at both t_2 and t_3 , Alice has a t_3 -obligation to $\langle K \rangle$ rather than $\langle L \rangle$ at t_4 , and, unlike $\langle L \rangle$, the content of $\langle K \rangle$ is shared by the content of $\langle C \rangle$. The shared content between $\langle C \rangle$ and $\langle K \rangle$ is exactly what we would expect *if*, contrary to what securitists claim at this point in the dialectic, Alice had a non-dependent t_1 -obligation to $\langle C, F, K \rangle$.

Goldman's version of securitism (1978: 194-195) which, as discussed in the previous section is, strictly speaking, not a version of maximalism, appears to rule out this response since, in *WI5*, although the act-set $\langle C, F, L \rangle$ is securable for Alice at t_1 in virtue of the fact that the performance of $\langle C \rangle$ would result in $\langle C, F, L \rangle$, $\langle C \rangle$ would *not* be causally efficacious if performed in the sense that part of the content of $\langle C \rangle$ (doing research for a paper at t_4) would *not* be actualized if Alice were to $\langle C \rangle$ at t_2 . Moreover, although Portmore's (2011: 193-194; 2019: ch. 5, fn. 10) version of securitism is more relaxed insofar as an act-set only needs to be securable in virtue of an immediately performable decision whose content is quite general,³² it is unclear whether

³² Cf. Brown (2018: 764-766) and Gustafsson (2014: 587-588).

Portmore would want to allow for such a mismatch between the content of $\langle C \rangle$ and the content of $\langle L \rangle$.

Perhaps securitists would say that in *WI5* Alice has a non-dependent t_1 -obligation to $\langle C, F \rangle$ from t_2 - t_3 as opposed to $\langle C, F, K \rangle$ from t_2 - t_4 , thereby avoiding the mismatch between the content of $\langle C \rangle$ and the content of Alice's non-dependent t_1 -obligation. However, this move appears untenable because, according to securitism, the act-set $\langle C, F \rangle$ is obligatory for Alice at t_1 at least partly because of its value, and its value is at least partly determined by the fact that if Alice were to $\langle C, F \rangle$ from t_2 - t_3 then she would $\langle L \rangle$ at t_4 . To illustrate this point further, suppose that the value of $\langle L \rangle$ is -100 as opposed to +5. In that case, according to securitism, Alice would have a non-dependent t_1 -obligation to $\langle B, D, I \rangle$ from t_2 - t_4 because this would be the best maximal act-set that is securable for Alice at t_1 .

A final way in which securitists might respond to this alleged dilemma is by simply denying that Alice has a t_1 -obligation to do anything at t_2 , and that the lack of an obligation is explained at least in part by Alice's special knowledge of the fact that, no matter what she does up until t_3 , she will lack a t_3 -obligation to $\langle L \rangle$ at t_4 . We think that a solution to the A/P debate should be able to handle the stipulations in *WI5* while affirming that Alice has a t_1 -obligation to do something at t_2 . For instance, possibilists have no difficulties maintaining that in *WI5* Alice has a t_1 -obligation to $\langle B \rangle$ at t_2 , and no knowledge of the kind that Alice has in *WI5* threatens this verdict. For example, if we further stipulate that in *WI5* Alice knows that, no matter what she does up until t_3 , she lacks a t_3 -obligation to $\langle I \rangle$ at t_4 , this is no obstacle to the possibilist's contention that Alice has a t_1 -obligation to $\langle B \rangle$ at t_2 .

We don't declare possibilism victorious, however, since it is subject to the worst outcome objection. Instead, we propose that an alternative view, hybridism, avoids the nonratifiability problem as well as the other problems that plague actualism and possibilism.

8. Hybridism

The preceding discussion suggests that important *desiderata* for the correct view in the A/P debate include (i) accommodating the intuitively correct verdicts rendered by both actualism and possibilism, while avoiding the (ii) not demanding enough (iii) bad behavior (iv) and worst outcome objections and, perhaps most importantly, (v) the nonratifiability problem. For reasons already given, a view that can accommodate (i)-(v) should be particularly appealing to consequentialists. The good news for consequentialists is that there are such views, which we refer to collectively as hybridism. In this final section, we'll first provide a formulation of a particular version of hybridism and then explain how hybridism captures the aforementioned *desiderata*. These considerations provide good reason for consequentialists to accept hybridism over its competitors.

Hybrid views are unique, in part, because they posit two distinct moral "oughts", one actualist in nature and one possibilist in nature. These oughts are meant to jointly track the insights of both actualism and possibilism, yet be immune from the five aforementioned objections. Given space limitations, we cannot provide a complete defense of any particular hybrid view here, so our goal is merely to make a *prima facie* case for hybridism by explaining how it satisfies the aforementioned *desiderata*. We will focus on one version of hybridism known as *single obligation hybridism (SOH)*. In its simplest form, *SOH* posits a possibilist moral *obligation* that picks out the criterion of right, which allows *SOH* to (a) accommodate the intuitively correct possibilist verdicts

and to (b) avoid the not demanding enough objection (c) the bad behavior objection and (d) the nonratifiability problem. In order to (a) accommodate the intuitively correct actualist verdicts and (b) avoid the worst outcome objection, *SOH* also posits an actualist moral *ought* that functions as a sort of decision procedure. In certain cases, it prescribes performing a wrong act at one time in order to avoid performing an even worse act at a different time. This moral ought is an action-guiding practical ought, not a moral obligation. *SOH* may be formulated more precisely as follows (Timmerman & Cohen, 2016: 682-683).

Single Obligation Hybridism:

Possibilist Moral Obligation: At t an agent S has a possibilist moral obligation to ϕ at 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 S can possibly perform an act.

Actualist Practical Ought: At t an agent S has most practical reason to ϕ at t' iff ϕ -ing at t' is under S 's control at t and ϕ -ing at t' is either (i) identical to the maximally-specific possibilist obligation that S has at t , (ii) a rationally permissible supererogatory act, or (iii) is the least rationally impermissible, all things considered, act-set presently under S 's control at t . There is an act-set that satisfies (iii) iff no act-set presently under S 's control at t satisfies conditions (i) or (ii).

At t an agent S has a dependent possibilist obligation to perform an act-set

Ψ iff (and because) at t S has a non-dependent possibilist moral obligation to perform an act-set Ψ^* , such that Ψ is a proper subset of Ψ^* . Additionally, at t , S non-dependently has most practical reason to perform act-set Ψ iff (and because) at t S non-dependently has most practical reason to perform an act-set Ψ^* , such that Ψ is a proper subset of Ψ^* .

The technical details of *SOH* are not centrally important for the purposes of this chapter. The most important elements of the view may be understood by considering the non-technical description of the view that preceded the formal definition and by considering *SOH*'s applications in particular cases. Consider *WII* once more. *SOH* entails that Alice has a possibilist moral obligation to $\langle B, D, H \rangle$ because this is the best maximal act-set that Alice can perform over time. This is the intuitively correct verdict. This feature allows *SOH* to avoid both the not demanding enough and bad behavior objections for obvious reasons. It also allows *SOH* to avoid the nonratifiability problem because Alice will always be obligated to perform each basic act that is part of the best set of acts she can perform at any given time. This possibilist feature of *SOH* makes it logically impossible for Alice to, at one time, have an obligation to perform an act at some future time and then, after only fulfilling her obligations, cease to have the future obligation in question. There is thus no reason to worry about *SOH* succumbing to the nonratifiability problem.

At the same time, *SOH* is immune from the worst outcome objection because the actualist ought, not the possibilist obligation, is action-guiding. The actualist ought prescribes performing the act that would result in the best outcome *from among the set of acts presently under the agent's control*. This practical ought, then, serves the purpose of minimizing wrongdoing in light of one's present circumstances. So, in *WII*, *SOH* entails that Alice practically ought to $\langle C \rangle$, which would

result in her performing $\langle C, F, K \rangle$. This captures the intuitively correct actualist verdict. *SOH* has this implication because, of the act-sets at t_2 under Alice's control at t_1 , performing $\langle C \rangle$ would result in the best outcome. So, *SOH* entails that Alice ought to perform a wrong act now (i.e. $\langle C \rangle$) in order to prevent herself from performing an even worse act later (i.e. $\langle I \rangle$). There is much more to be written in favor of (and against) hybridism. Indeed, there is much more to be written about each of these views. The considerations we raise in this section, however, should provide a least a presumptive case in favor of hybridism.

9. Conclusion

This chapter served a few related goals. We first traced the origins of the A/P debate to a debate in the deontic logic literature about the coherence of utilitarianism. We then discussed the relationship between time, an agent's abilities, and their moral obligations before introducing a precise version of the wedding invitation case that would guide us through the chapter. Using that case, we reviewed the four primary views in the A/P debate, explaining their benefits and drawbacks from a consequentialist perspective. In doing so, we attempted to further the A/P debate by arguing that all forms of securitism (including Ross's *MWSS*) are subject to the nonratifiability problem. Moreover, we argued that the nonratifiability problem is more serious than it may initially seem, as it either violates OIC or is committed to an implausible position on the relationship between dependent and non-dependent obligations. We ended by making a positive case for hybridism which, we argued, avoids the nonratifiability problem in addition to each of the other problems discussed in the chapter.³³

³³ We are grateful to Doug Portmore for very helpful feedback on this chapter. This chapter is the product of full and equal collaboration between its authors.

References

- Almeida, M. (1992). The paradoxes of Feldman's neo-utilitarianism. *Australasian Journal of Philosophy* 70(4), 455–468.
- Bergström, L. (1968). Utilitarianism and deontic logic. *Analysis* 29(2), 43–44.
- Brown, C. (2018). Maximalism and the structure of acts. *Noûs* 52(4), 752–771.
- Bykvist, K. (2002). Alternative actions and the spirit of consequentialism. *Philosophical Studies* 107(1), 45–68.
- Cariani, F. (2016). Consequence and contrast in deontic semantics. *Journal of Philosophy* 113(8), 396–416.
- Carlson, E. (1995). *Consequentialism reconsidered*. Dordrecht: Kluwer.
- Castañeda, H. (1968). A problem for utilitarianism. *Analysis* 28(4), 141–142.
- Cohen, Y., & Timmerman, T. (2016). Actualism has control issues. *Journal of Ethics and Social Philosophy* 10(3), 1–19.
- Dorsey, D. (2009). Headaches, lives and value. *Utilitas* 21(1), 36–58.
- Dreier, J. (2011). In defense of consequentializing, in M. Timmons (ed.), *Oxford studies in normative ethics, volume 1*, 97–119. New York: Oxford University Press.
- Feldman, F. (1986). *Doing the best we can: An essay in informal deontic logic*. Boston: D. Reidel.

Goldman, H.S. (1976). Dated rightness and moral imperfection. *The Philosophical Review* 85(4), 449–487. [Now Holly M. Smith]

Goldman, H.S. (1978). Doing the best one can, in A.I. Goldman and J. Kim (eds.), *Value and morals: Essays in honor of William Frankena, Charles Stevenson, and Richard Brandt*, Dordrecht: D. Reidel: 185–214. [Now Holly M. Smith]

Greenspan, P.S. (1978). Oughts and determinism: A response to Goldman. *The Philosophical Review* 87(1), 77–83.

Gustafsson, J.E. (2014). Combinative consequentialism and the problem of act versions. *Philosophical Studies* 167(3), 585–596.

Jackson, F. (1985). On the semantics and logic of obligation. *Mind* 94(374), 177–195.

Jackson, F. (2014). Procrastinate revisited. *Pacific Philosophical Quarterly* 95(4):634–647.

Jackson, F., & Pargetter, R. (1986). Ought, options, and actualism. *The Philosophical Review* 95(2), 233–255.

Kiesewetter, B. (2018). Contrary-to-duty scenarios, deontic dilemmas, and transmission principles. *Ethics* 129(1), 98–115.

Lewis, D. (1973). *Counterfactuals*, Oxford: Blackwell Publishers and Cambridge, MA: Harvard University Press, reprinted with revisions, 1986.

Norcross, A. (1998). Great harms from small benefits grow: How death can be outweighed by headaches. *Analysis* 58(2), 152–158.

- Pollock, J.L. (1976). *Subjunctive reasoning*. Dordrecht: D. Reidel.
- Portmore, D.W. (2009). Consequentializing. *Philosophy Compass* 4(2), 329–347.
- Portmore, D.W. (2011). *Commonsense consequentialism*. New York: Oxford University Press.
- Portmore, D.W. (2017). Maximalism versus omnism about permissibility. *Pacific Philosophical Quarterly*. 98(S1): 427–452.
- Portmore, D.W. (2019). *Opting for the best: Oughts and options*. New York: Oxford University Press.
- Prawitz, D. (1968). A discussion note on utilitarianism. *Theoria* 34(1), 76–84.
- Prawitz, D. (1970). The alternatives to action. *Theoria* 36(2), 116–126.
- Ross, J. (2012). Actualism, possibilism, and beyond, in M. Timmons (ed.) *Oxford studies in normative ethics, volume 2*, 74–96. New York: Oxford University Press.
- Scanlon, T. (1998). *What we owe to each other*. Cambridge, Mass.: Belknap Press.
- Sobel, J.H. (1976). Utilitarianism and past and future mistakes. *Noûs* 10(2), 195–219.
- Stalnaker, R. (1968). A Theory of conditionals, in N. Rescher (ed.), *Studies in logical theory* (American philosophical quarterly monograph series: volume 2), 98–112. Oxford: Blackwell.
- 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), 1509–1528.

Timmerman, T. (in press). Effective altruism's underspecification problem, in H. Greaves and T. Pummer (eds.) *Effective altruism: Philosophical issues*. New York: Oxford University Press.

Timmerman, T., & Cohen, Y. (2016). Moral obligations: actualist, possibilist, or hybridist? *Australasian Journal of Philosophy* 94(2), 672–686.

Timmerman, T., & Cohen, Y. (2019). Actualism and possibilism in ethics, in E.N. Zalta (ed.), *The Stanford encyclopedia of philosophy* (summer 2019 edition), URL = <<https://plato.stanford.edu/archives/sum2019/entries/actualism-possibilism-ethics/>>.

Timmerman, T., & Swenson, P. (2019). How to be an actualist and blame people, in D. Shoemaker (ed.) *Oxford studies in agency and responsibility, volume 6*. New York: Oxford University Press.

Vessel, J. (2009). Defending a possibilist insight in consequentialist thought. *Philosophical Studies* 142(4), 183–195.

Wedgwood, R. (2009). Against actualism, PEA Soup, posted on September 11, 2009. Accessed on April 3rd, 2019 from <http://peasoup.typepad.com/peasoup/2009/09/against-actualism.html>

Westphal, F. (1972). Utilitarianism and "conjunctive acts": A reply to professor Castañeda. *Analysis* 32(3), 82–85.

Woodard, C. (2009). What's wrong with possibilism. *Analysis* 69(2), 219–226.

Zimmerman, M.J. (1986). Subsidiary obligation. *Philosophical Studies* 50(1), 65–75.

Zimmerman, M.J. (1990). Where did I go wrong? *Philosophical Studies* 59(1), 55–77.

Zimmerman, M.J. (1996). *The concept of moral obligation*. Cambridge: Cambridge University Press.

Zimmerman, M.J. (2006). The relevant risks to wrongdoing, in K. McDaniel, J.R. Raibley, R. Feldman, and M.J. Zimmerman (eds.), *The good, the right, life and death: Essays in honor of Fred Feldman*, 151–72. New York: Ashgate:

Zimmerman, M.J. (2017). Prospective possibilism. *Journal of Ethics* 21(2), 117–150.