Fault Lines in Ethical Theory

Shyam Nair *

August 2019

Abstract

The verdicts standard consequentialism gives about what we are obligated to do crucially depend on what theory of value the consequentialist accepts. This makes it hard to say what separates standard consequentialist theories from non-consequentialist theories. This article discusses how we can draw sharp lines separating standard consequentialist theories from other theories and what assumptions about goodness we must make in order to draw these lines. The discussion touches on cases of deontic constraints, cases of deontic options, and cases involved in the so-called "actualism"/"possibilism" debate. What emerges is that there are various interesting patterns relating the different commitments of consequentialism, different principles about obligation and about goodness, and different rules concerning how facts about values determine facts about obligation.

KEYWORDS: constraints, options, actualism, possibilism, consequentializing, axiology, deontic logic, decision rules

Introduction

When we introduce students to moral philosophy, we draw sharp contrasts between different ethical theories. Chief among these contrasts is the contrast between consequentialism and non-consequentialist moral theories. We often illustrate the contrast with cases like the Williams's integrity challenges (Williams 1973) and Foot's Trolley Problem (Foot 1967 [1979]).

But the contemporary literature about moral philosophy does not fit well with this tidy picture. As the articles in this handbook attest, consequentialism is a label for a startling variety of different theories that can often make different predictions about these cases. Indeed, some have gone so far as to argue that consequentialism is such an accommodating framework that it can deliver the same verdicts as any plausible non-consequentialist theory.[1] This is because, these theorists argue, there is always

---

[1]Important proponents of the approach include Dreier 1993, Dreier 2011, Louise 2004, Oddie and Milne 1991, Portmore 2007, Portmore 2009. Important criticisms include Brown 2011, Schroeder 2007. An important recent work that is relevant to this discussion but harder to categorize is Dietrich and List 2017.

some ranking of outcomes according to their goodness that can deliver whatever result about a case that we wish.

This article investigates to what extent we can draw lines that separate standard consequentialist theories from other theories. In my usage, a fault line is a case or really a structural description of a class of cases that can be used to separate moral theories because the theories give different verdicts about the class of cases. I will be interested in what fault lines there are and in the closely related question of how to establish that these fault lines exist.

We begin by exploring two of the most studied kinds of cases that are supposed to separate consequentialist theories from deontological theories, cases of deontic constraints and cases of deontic options (§1). We then discuss another family of fault lines that can be described using certain simple abstract principles (§2). These fault lines separate standard consequentialist theories from certain non-consequentialist theories that nonetheless give the value of outcomes an important place in determining the deontic status of acts (§3). Finally, we close by commenting on how our discussion bears on broader issues in ethical theory (§4).

1 Deontic Constraints and Options

Cases of deontic constraint come in a variety of forms.[2] Here is a simple case:

> Mac faces a choice of whether or not to murder an innocent person, Jack. If Mac does not murder Jack, Frank and Hank will each murder a different innocent person. If Mac does murder Jack, he will prevent Frank and Hank from murdering.

Many believe that it would be wrong for Mac to murder Jack even though Mac's murdering Jack would prevent Frank and Hank from murdering two innocent people. Many non-consequentialist theories give this result about the case. On the other hand, it is hard to see how standard act consequentialism——$\mathcal{S}$ is obligated to do x iff the outcome of $\mathcal{S}$'s doing x is better than the outcome of $\mathcal{S}$'s refraining from doing x—can get this result.[3] Given that the outcome of Mac's murdering Jack contains just one innocent person being murder and given that the outcome of Mac's refraining from murdering Jack contain two innocent murders, it seems that the outcome of Mac's murdering Jack is better than the outcome of Mac's refraining from murdering Jack. Thus, standard act consequentialism appears to be committed to the idea that Mac is obligated to murder.

And cases like Mac's are just one of a class of structurally similar cases that raise similar problems such as the following kinds of cases:

- cases where an agent can break a promise in order to prevent others from breaking their promises
- cases where an agent can forgo providing a benefit for one of her nearest-and-dearest in order to prevent others from forgoing providing a benefit for one of their nearest-and-dearest

---

[2]These cases are sometimes called cases of agent-centered restrictions and are taken to show that consequentialism only allows agent-neutral considerations to be relevant. The classic discussion in this vein is perhaps Scheffler 2003 [1982] (see also Darwall 1986 which, among other things, provides a useful historical and conceptual background for understanding these cases).

[3]I stipulate that $\mathcal{S}$'s refraining from doing x is $\mathcal{S}$ not doing x.

One standard gloss on what the structural similarity between these cases are is that they are cases where an agent faces a choice of whether to perform an act where performing that act will prevent others from performing an act of the same morally relevant type. According to some non-consequentialist theories, there are cases where it is wrong for an agent to perform an act (e.g., murder, break a promise) even though doing so would prevent others from performing an act of the same morally relevant type (e.g., murder, break a promise). But according to standard consequentialism it appears there can be no such cases.

Cases of deontic options, like cases of deontic constraints, come in a variety of forms. Here is a simple case of this sort:

> Carol works for a non-profit helping the global poor. She works directly with those affected by poverty and is passionate about it. She could quit her job and work in finance and donate all her money to charity. She would help slightly more people by donating money rather than working directly with the global poor.

Many people believe that it is permissible for Carol to continue working for the non-profit and permissible for her to work in finance. Standard consequentialism, on the other hand, appears to entail that Carol is required to quit her job and work in finance because this will help slightly more people. Many non-consequentialist theories recognize a special dispensation permitting an agent to pursue the projects that matter the most to her, but standard consequentialism is accused of being incompatible with this.

1.1 Dimensions of Importance

Suppose that it has been shown that consequentialist and non-consequentialist theories are committed to giving different results about cases of deontic constraints and cases of deontic options. These cases then would be fault lines that separate these theories. Can anything more be said about these fault lines?

One way to think about these questions is to ask whether the fault lines target only some of the core commitments of consequentialism or all of them. There are at least three core commitments of consequentialism. The first two are the commitments we get from separating the biconditional that is the statement of consequentialism into a pair of conditionals:

> LEFT-TO-RIGHT: if S is obligated to do x, then the outcome of S's doing x is better than the outcome of S's failing to do x
>
> RIGHT-TO-LEFT: if the outcome of S's doing x is better than the outcome of S's failing to do x, then S is obligated to do x

The third commitment is an entailment of consequentialism that is sometimes called the "compelling idea", the idea that it is always permissible to do what is best:

> COMPELLING IDEA: if the outcome of S's doing x is the best outcome, it is permissible for S to do x

We can now ask the more specific question of what cases of constraints and options say about each of these theses.

Begin with cases of constraints. If it is true that the outcome of Mac's murder is better than the outcome of Mac's refraining from murder and it is true that Mac is not obligated to murder and that Mac is obligated to refrain from murder, then cases of constraints are incompatible with all three theses. To see it is incompatible with

LEFT-TO-RIGHT consider that Mac is obligated to refrain from murdering even though the outcome of murdering is better than the outcome of refraining from murdering. To see that it is incompatible with RIGHT-TO-LEFT and COMPELLING IDEA, consider that the outcome of Mac murdering is best outcome available to Mac and yet it is not permissible for Mac to murder. On the other hand, cases of deontic options are only incompatible with RIGHT-TO-LEFT. In cases of deontic option it is merely permissible (so not obligatory) for Carol to work in finance even though the outcome of her doing so is better than the outcome of her not doing so (and dedicating her life to working at the non-profit). On the other hand, these cases are no threat to RIGHT-TO-LEFT because no act is obligatory in this case. And these cases pose no threat to Compelling Idea because the act with the best outcome is permissible.

Thus, there is a sense in which cases of deontic constraints are more *deeply* in tension with the consequentialist framework than cases of deontic options. I will, then, say we can analyze putative fault lines according to how deeply incompatible with consequentialism they are where one measure of depth is which commitments of consequentialism must be rejected.

Of course, unlike the ordinary notion of depth, this notion cannot be numerically measured. Instead, it is better understood as a useful indicator. So though we may not be able to so precisely compare the depth of fault lines, we can still use adjectives related to depth as helpful labels that tell us to look more closely at exactly which commitments are rejected by each fault line.

1.2 Assumptions about Goodness

In order to establish that consequentialism is incompatible with cases of deontic constraints and cases of deontic options, we need to make substantive assumptions about the goodness of outcomes, e.g., that the outcome containing two murders was worse than the outcome containing one.

And these assumptions are initially quite plausible. But what is to prevent the consequentialist from rejecting these assumptions in favor of a different account of goodness? For example, what is to prevent the consequentialist from claiming that keeping one's hands clean from murder is better than preventing those who already have murderous consciences from murdering?

More generally, consequentialist who are pluralists about the good have considerable resources to accommodate putative fault lines. Campbell Brown observes that to those who claim that a certain case is incompatible with consequentialism:

> [...] the consequentialist [...may] reply. "Your story presupposes a certain account of what makes consequences better or worse, a certain 'theory of the good,' as we consequentialists like to say. Consequentialism, however, is not wedded to any such theory. We already knew that combining consequentialism with some theories of the good would have implausible results; that's what utilitarianism has taught us. In order to reconcile consequentialism with the view that this action you've described is wrong, we need only to find an appropriate theory of the good, one according to which the consequences of this action would not be best. You say you're concerned about the guy's rights? No worries; we'll just build that into your theory of the good. Then you can be a consequentialist too." (Brown 2011: 749-750)

4

Of course, there may be arguments against the theory of goodness that consequentialism requires to get the desired results in these cases. But there is no easy road from simple verdicts about cases to the rejection of consequentialism. Additional argument concerning what is good is needed.

If the strategy that Brown outlines works, it would show that we cannot establish the existence of fault lines separating consequentialist from non-consequentialist theories while being neutral about what theory of goodness is correct. I call the question of what it takes to establish a fault line the question of on *what grounds* the fault line can be established. And what we have just seen is that there is an interesting question of whether any fault line can be established on *neutral* grounds.

This gives us a second dimension of evaluation by which to consider various fault lines.

1.3   There are No Deep Fault Lines on Neutral Grounds

We have already seen that given certain assumptions about goodness, we can establish that cases of deontic constraints form a deep fault line. But we have also seen that these assumptions are substantive. In this subsection, we will see that though some fault lines *can* be established on neutral grounds, no *deep* fault line can be established on *neutral* grounds.

1.3.1   A Fault Line on Neutral Grounds

We begin by seeing that we can establish that cases of deontic constraints are incompatible with consequentialism on *neutral* grounds. The only assumptions about goodness that we will make are standard logical assumptions about the betterness relation (e.g., that it is irreflexive, transitive, and complete).[4]

Consider a case where if Mary doesn't murder, John will and where if John doesn't murder, Mary will. Intuitively, this is a case of constraints where it is wrong for Mary to murder and wrong for John to murder.

But consider that there are only two possible ways things could go:

w_1: Mary murders, John doesn't

w_2: Mary doesn't murder, John does

No matter how these outcomes are ordered with respect to betterness, we must reject either the claim that it is obligatory for Mary to not murder or the claim that is is obligatory John to not murder (and not obligatory for them to murder). For suppose w_1 is better than w_2, in this case consequentialism entails that it obligatory for Mary to murder. Analogously, if w_2 is better than w_1, it follows that it is obligatory for John to murder. Finally, if w_1 and w_2 are equally good, it follows that it is not obligatory for either of them to refrain murder. Thus, it follows that no matter what theory of goodness we accept consequentialism is incompatible with cases of deontic constraint.

1.3.2   Impossibility and Modesty

That said, this argument does not establish that there is a deep fault line for it is compatible with Left-to-Right and Compelling Idea. If w_1 and w_2 are equally ranked, $Left-to-Right$ must hold because no act is obligatory in such a case and $CompellingIdea$ must hold because all acts are permissible in such a case. More

---

[4]This neutral argument is essentially the one in Brown 2011: 761-3.

generally, in order to show that either of these theses fails, one outcome must be ranked ahead of another. This shows that no neutral argument can ever establish a deep fault line because there is nothing about the logical structure of goodness that entails that one outcome is strictly better than another.

This teaches us that in order to establish deep fault lines, we will need to appeal to some substantive principle that tells us one outcome is better than another. But these principles need not be as strong as the claim that violating rights is bad or donating to charity is good. Instead, we can rely on *modest* theoretical principles.

One such principle is a Pareto-like principle that I will call UNANIMITY

> UNANIMITY: if every agent ought to prefer an outcome w_i to another outcome w_j, then w_i is better than w_j

This is a high-level minimal principle connecting the preferences that every agent ought to have with goodness.[5] Though I do not believe that there is any snappy argument in favor of this principle, it is quite plausible.

In other work (Nair 2014), I argue that if we accept this principle, we can show that a certain class of cases of deontic constraints form a deep fault line. Since the principle is neither a purely logical principle nor a more substantive claim about goodness, I say we can establish that these cases form a deep fault line on relatively modest (albeit not neutral) grounds. The argument for this is too lengthy to be rehearsed here.[6] And the details are in any case inessential. But what matters for our purposes is that we may uncover certain modest principles about value that can be used to establish fault lines in ethical theory. The next section adopts this approach.

2 Structural Descriptions of Cases and Standard Consequentialism

As we have seen, typically fault lines are paradigmatic examples (e.g., cases where common sense morality suggests it is wrong to murder even to prevent more murders) or specified with informal glosses (e.g., cases where it is wrong to do an act even though that act would prevents the performance of more acts of the same morally relevant type). But little work has been done to describe these fault lines formally.

This is understandable as these cases have quite complex normative and causal structures. Nonetheless, in this section, we will consider how simple formal principles can specify fault lines.

2.1 The Logical Structure of Cases

We focus on extremely simple cases involving only what an agent at a time is obligated to do (though these acts may occur at distinct times). We will use O as an obligation operator take it to be implicitly indexed to an agent at a time. We will also assume that

---

[5]Though it abuses language, I treat 'ought' and 'obligatory' and their variants as synonymous throughout this paper.

[6]Roughly, the argument proceeds by claiming that cases of deontic constraints only arise when the preferences that agents ought to have differ. A subclass of such cases will be cases with the structure of a prisoner's dilemma. In this subclass, the UNANIMITY principle forces the "cooperate-cooperate" outcome to be ranked ahead of the "defect-defect" outcome in such a way that allows us to show that consequentialism is deeply incompatible with cases of deontic constraints of this structure.

O officially is a propositional operator rather than an operator on acts themselves. But often, when speaking informally, we will freely switch between treating the prejacents of O as propositions or as acts. In this setting, we can write down some simple sentences and take them to characterize a class of cases. For example, presumably the sentence '$O(A)$ or $\neg O(A)$' characterizes every case. And presumably every theory says 'yes' to the existence of this class and 'no' to the existence to a non-empty complement of this class. So it does not form a fault line.

But there are other more interesting cases as well. Here is one example:

> AGGLOMERATION: $O(A) \land O(B) \rightarrow O(A \land B)$

AGGLOMERATION is a structural description of a class of cases in which either $O(A) \land O(B)$ is false or $O(A \land B)$ is true. The complement of this class is the cases where $O(A) \land O(B)$ is true and $O(A \land B)$ is false. For AGGLOMERATION to specify a fault line would be for there to be a dispute among theories about whether the class specified by it or by its complement are non-empty. As we will see, there is such a dispute about whether the complement of this class of case is non-empty.

In a similar vein, we can ask about the class of cases and complement of the class of cases characterized by the following principles:

> INHERITANCE: If A entails B, then $O(A)$ entails $O(B)$
>
> NO CONFLICTS: If $O(A_1)$ is true, $O(A_2)$ is true, ..., and $O(A_n)$ is true, then $\{A_1, A_2, \ldots, A_n\}$ is consistent
>
> NO STRICT CONFLICTS:$\neg(O(A) \land O(\neg A))$

Later, we will discuss the plausibility of these principles (§3.5). But for now, we turn to an initial discussion of how consequentialism as standardly formulated is incompatible with these principles.

2.2 Standard Consequentialism and the Ubiquity of Fault Lines

As discussed early, standard act consequentialism is the following claim:

> $\mathcal{S}$ is obligated to do x iff the outcome of $\mathcal{S}$'s doing x is better than the outcome of $\mathcal{S}$'s failing to do x

where we say for a possible world, w:

> w is the outcome of $\mathcal{S}$'s doing(/refraining from doing) x iff if $\mathcal{S}$ were to do(/refrain from doing) x, then w would obtain[7]

We will look in greater detail at the features of this formulation in §3.1. But for now, what I wish to point out is that AGGLOMERATION, INHERITANCE, and NO CONFLICTS appear to be fault lines in that standard consequentialism accepts the existence of cases that falsify these principles while other theories accept these principles so are committed to there being no such cases.

Many cases in the literature have been offered that attest to standard consequentialism's rejection of these principles.[8] Here we can consider one due to Michael Zimmerman:

---

[7]I stipulate '$\mathcal{S}$ refrains from doing x' is equivalent to the sentential negation of '$\mathcal{S}$ does x'.

[8]Cases of this sort are discussed in the so-called "actualism/possibilism" debate. See Timmerman and Cohen, this volume (and the citations therein) for a rich discussion of this literature. There are many interesting relations between their discussion and what is to follow. Unfortunately, I am unable to pursue these at this time.

I have been invited to attend a wedding. The bride-to-be is a former girlfriend of mine; it was she who did the dumping. Everyone, including me in my better moments, recognizes that she was quite right to end our relationship; we were not well suited for one another, and the prospects were bleak. Her present situation is very different; she and her fiancé sparkle in one another's company, spreading joy wherever they go. This irks me to no end, and I tend to behave badly whenever I see them together. I ought not to misbehave, of course, and I know this; I could easily do otherwise, but I do not. The wedding will be an opportunity for me to put this boorishness behind me, to grow up and move on. The best thing for me to do would be to accept the invitation, show up on the day in question, and behave myself. The worst thing would be to show up and misbehave; better would be to decline the invitation and not show up at all. (Zimmerman 2006: 153)

Zimmerman adds to fill out the case "if I accepted the invitation, I would show up and misbehave (whereas I would not do this if I declined). I need not misbehave (for, as noted, I could easily do otherwise); nonetheless, this is what I would in fact do" (ibid.: 153).

Here Zimmerman can accept the invitation in two ways. One way would be to accept the invitation and go on to behave well; the other to accept the invitation and go on to behave poorly. If he were to accept the invitation, he would go on to behave poorly. He can, on the other hand, decline the invitation. The best outcome is the outcome in which Zimmerman accepts and behaves well. The middle outcome is the outcome in which Zimmerman declines the invitation. The worst outcome is the outcome in which Zimmerman accepts and behaves poorly. The outcome of accepting is the outcome in which Zimmerman accepts and behaves poorly.

According to standard consequentialism, Zimmerman is obligated to decline because the outcome of Zimmerman declining is the middle outcome which is better than the outcome of Zimmerman not declining (i.e., accepting) which is the worst outcome. By similar reasoning, it is not the case that Zimmerman is obligated to accept. But Zimmerman is obligated to accept and behave well because the outcome of this is best.

If we let '$Accept$' expressed the proposition that Zimmerman accepts the invitation, '$Behave\ Well$' express the proposition that Zimmerman behaves well, and '$Decline$' express the proposition that Zimmerman declines the invitation, we have $O(Accept \wedge Behave\ Well)$, $O(Decline)$, $\neg O(Accept)$. This falsifies INHERITANCE. Next, since one cannot accept, behave well, and decline, this also falsifies NO CONFLICTS.

Finally, we assume throughout an impossible claim is never obligatory.[9] So we have $\neg O(Accept \wedge Behave\ Well \wedge Decline)$. This, then, falsifies AGGLOMERATION as well. Thus these principles appear to represent a fault lines in ethical theory in that standard consequentialism cannot accept them and other theories do accept them.[10]

But do examples like Zimmerman's rest on substantive assumptions about goodness?

---

[9]This can be justified by the Kantian idea that one is obligated to do something only if one can do it or by the assumption that an impossible situation is worse than any possible situation.

[10]NO STRICT CONFLICTS is something that standard consequentialism accepts. This is because A cannot have an outcome that is better than the outcome of $\neg A$ while $\neg A$ has a better outcome than the outcome of A at the same time.

2.3   Non-Neutral Grounds

To check whether they do, let's begin by understanding the structure of Zimmerman's example: There is an act A (e.g., *Accept*) that can be done in two incompatible ways, $A \wedge B$ (e.g., *Accept $\wedge$ Behave Well*) and $A \wedge \neg B$ (e.g., *Accept $\wedge$ $\neg$Behave Well*).[11] These acts result in distinct incompatible outcomes, $w_{(A \wedge B)}$ and $w_{(A \wedge \neg B)}$. Further, we assumed that there is another act C (e.g., *Decline*) that has a third distinct outcome, w_C. Finally, we claimed that if one were to do A (e.g., *Accept*), one would do it some particular way such as $A \wedge \neg B$ (e.g, *Accept $\wedge$ $\neg$Behave Well*). Evidently, these claims make no assumptions whatsoever about the goodness or badness of outcomes. And they are, in any case, thoroughly innocuous.

But the case relied on some further assumptions. In particular, I claimed that we can rank the outcomes so that w_C is strictly in the middle (i.e., it is strictly better than exactly one of $w_{(A \wedge B)}$ and $w_{(A \wedge \neg B)}$ and strictly worse than exactly one of $w_{(A \wedge B)}$ and $w_{(A \wedge \neg B)}$). In the example, I claimed the outcome of declining the invitation is strictly better than the outcome of accepting and behaving poorly and strictly worse than the outcome of accepting and behaving well.

But we already know for the reasons given in §1.3.2 that these assumptions about goodness must be non-neutral: The assumptions entail that some outcome is ranks ahead of another outcome. But there is nothing about the minimal structural properties of goodness that would tell us this. So the argument that standard consequentialism must reject the above principles relies on some substantive assumptions about goodness.

2.4   Modest Grounds

But what are the weakest assumptions that we actually need to make about goodness in order to establish this fault line? Or to frame this issue in a different way, is there a theory of goodness that can we supplement standard consequentialism with so as to show that it can accept INHERITANCE, AGGLOMERATION, and NO CONFLICTS?

We already saw that in Zimmerman's example the act of accepting can be done in two ways, one can accept and go on to behave well or accept and go on to behave poorly. And in Zimmerman's example there was also the act of declining which is incompatible with all of these acts. The crucial additional value assumption that is made in Zimmerman's example is that the value of the outcome of declining can be strictly in between the value of these two different ways of accepting. It assumes that we can in fact pry apart the value of two outcomes in which one accepts.

This suggests that if we adopt a theory of goodness that does not allow us to pry apart the value of two outcomes in which an agent accepts, Zimmerman's case would not be enough to establish that consequentialism is incompatible with INHERITANCE, AGGLOMERATION, and NO CONFLICTS. More generally, it suggests the conjecture that if value of the outcomes in which a given act, A, occurs cannot be "splintered" in the sense that there is an incompatible act B whose outcome is strictly in two distinct outcomes in which A occurs, then standard consequentialism can accept INHERITANCE, AGGLOMERATION, and NO CONFLICTS. This conjecture can be formalized as follows:

> VALUE NON-SPLINTERING: if A is true at w_i and at w_j, then for any w_k where A is false, if w_i is strictly better than (/worse than/equally good as) w_k, then w_j is better than (/worse than/equally good as) w_k

---

[11]Others such as Brown 2018and Portmore 2019: ch. 4 have also thought of this case (or cases similar to it) as a case where there are multiple versions or ways of doing an act.

where A is a statement that the agent does some act (e.g, $\mathcal{S}$ does x) and where w_i, w_j, w_k are outcomes.[12] According to this principle, It cannot be that the outcome of *Decline* is strictly in between the outcomes of *Accept* $\wedge$ *Behave Well* and the outcome of *Accept* $\wedge$ $\neg$*Behave Well*. This is because the outcome of *Accept* $\wedge$ *Behave Well* and the outcome of *Accept* $\wedge$ $\neg$*Behave Well* are both outcomes in which *Accept* holds. And the principle tells us that no outcome where *Accept* does not hold can come strictly between two outcome in which *Accept* holds. Thus, if this principle holds, we cannot use Zimmerman's case to establish that consequentialism is incompatible with INHERITANCE, AGGLOMERATION, and NO CONFLICTS.

Of course, it is unsurprising that we get these results in Zimmerman-type cases as the principle is tailored to handle that case. But our more general conjecture is also true: Standard consequentialism paired with VALUE NON-SPLINTERING entails INHERITANCE, AGGLOMERATION, and NO CONFLICTS. Proposition 1-Proposition 3 in §A provide the relevant proofs of the conjecture.

VALUE NON-SPLINTERING and its negation are non-trivial claims. That said, they are quite different from the claim, e.g., that protection of rights makes no contribution to goodness. One obvious difference is its generality.

Another difference is that VALUE NON-SPLINTERING is quite implausible in its own right. One way to think about VALUE NON-SPLINTERING is that it says that for a given act A the value of A has a kind of lexical priority in ordering outcomes in which A with respect outcomes in which $\neg A$. Outcomes in which A may be better or worse than or equal to one another. But when it comes to comparisons to outcomes in which $\neg A$, it is only the presence of A that matters. While this property may be sensible for certain acts that are especially morally awful or especially morally good, it is not sensible for every act. Consider ordinary acts such as the act of deciding to spend time reading a book, going to the movies, or eating dinner. It is implausible that anything that one can do that results in these acts occurring ranks the same as any other outcome in which these acts occur. Surely, reading a very good book may be better than not reading any book which in turn is better than reading a very bad book. This is something that VALUE NON-SPLINTERING forbids.

For this reason, then, I believe the rejection of VALUE NON-SPLINTERING is a very modest commitment about goodness. And as such, it can be shown on modest grounds that each of the principles of AGGLOMERATION, INHERITANCE, and NO CONFLICTS each form a fault line separating standard consequentialism which must reject all of these claims from other theories which can accept these claims.

2.5 Depth

How deep is this fault line? First, the fault line does not falsify COMPELLING IDEA (which recall says that if the outcome of $\mathcal{S}$'s doing x is the best available outcome, then it is permissible for $\mathcal{S}$ to do x). This is because in these cases the act which has the best outcome (e.g. accepting and behaving) is obligatory and hence is permissible according to standard consequentialism and according to how cases like Zimmerman's are standardly presented.

Whether both LEFT-TO-RIGHT (i.e., if $\mathcal{S}$ is obligated to do x, then the outcome of $\mathcal{S}$'s doing x is better than the outcome of $\mathcal{S}$'s failing to do x) and RIGHT-TO-LEFT (i.e.,

---

[12] The qualification that 'w_i, w_j, w_k are outcomes' restricts our attention to possible worlds that would result from some act available to the agent. So the principle does not apply to possible worlds that could not result from some act available to the agent

if the outcome of $\mathcal{S}$'s doing x is better than the outcome of $\mathcal{S}$'s failing to do x, then $\mathcal{S}$ is obligated to do x) fail depends on exactly how VALUE NON-SPLINTERING fails. To illustrate, suppose in Zimmerman's case *Accept* $\wedge$ *Behave Well* has an outcome that is strictly better than the outcome of $\neg$*Accept* which in turn is strictly better than the outcome of *Accept*. This is an instance of the failure of VALUE NON-SPLINTERING that has the following structure:

> TOTAL VALUE SPLINTERING: there is an act A such that w_i and w_j are outcomes in which A occurs and w_k is an outcome in which $\neg A$ occurs and w_i is strictly better than w_k and w_k is strictly better than w_j

In cases, with this structure, both LEFT-TO-RIGHT and RIGHT-TO-LEFT fail.

To see this, suppose INHERITANCE, AGGLOMERATION, and NO CONFLICTS hold so that it is obligatory to accept and behave, that it is obligatory to accept, and that it is not obligatory to decline. Here LEFT-TO-RIGHT fails because it is obligatory to accept even though the outcome of accepting is not strictly better than the outcome of declining. And here RIGHT-TO-LEFT fails because the outcome of declining is strictly better than the outcome of accepting but it is not obligatory to decline.

On the other hand, VALUE NON-SPLINTERING can fail without TOTAL VALUE SPLINTERING holding. For example, suppose we hold that outcome of accepting and behaving is strictly better than the outcome of declining but then only claim that the outcome of accepting (and not behaving) is equally good as the outcome of declining. This is not an instance of TOTAL VALUE SPLINTERING because the outcome of accepting and of declining are equally good. Instead, this is an instance of the following general structure:

> PARTIAL VALUE SPLINTERING: there is an act A such that w_i and w_j are outcomes in which A occurs and w_k is an outcome in which $\neg A$ occurs and w_i is strictly better than w_k and w_k is equally good w_j

As the interested reader can verify, if we only have instance of PARTIAL VALUE SPLINTERING, LEFT-TO-RIGHT fails, but RIGHT-TO-LEFT does not fail. Since I am skeptical there are any plausible grounds for merely accepting PARTIAL VALUE SPLINTERING, I will not explore this more restricted failure of VALUE NON-SPLINTERING further. Instead, I conclude that on relatively modest grounds we have located a fault line that is a kind of intermediate between the depth of cases of deontic constraints and the shallowness of cases of deontic options.

3 Further Dimensions of Depth

There are further questions we can ask to assess the depth of a fault line. We can consider what kinds of theories are separated by a fault line and how interesting these theories are. And we can also consider what kinds cases constitute the fault line and how interesting these cases are. We spend most of this section considering the first issue (§3.2-3.4), but close with a brief discussion of the second (§3.5).

As we will see, the cases that illustrate the failures of the principles that we are discussing separate standard consequentialism from other kinds of value-based theories. In order to see this clearly, we begin by highlighting the main features of standard consequentialist.

3.1 Basic Resources

Recall that standard consequentialism is the following claim:

> $\mathcal{S}$ is obligated to do x iff the outcome of $\mathcal{S}$'s doing x is better than the outcome of $\mathcal{S}$'s failing to do x

where we say for a possible world, w:

> w is the outcome of $\mathcal{S}$'s doing x iff if $\mathcal{S}$ were to do x, then w would obtain

Let us look at the basic resources involved in this formulation. It, of course, involves explaining what is obligatory in terms of value. It also involves the notion of an outcome. As I have stipulatively defined it an outcome is a possible world (i.e., a maximally specific way things could be) and it is the possible world that would result if the agent did the act.

This notion of an outcome is in one way quite broad. It is broad in the sense that it is very inclusive: the outcome of an act is not merely its causal consequences. Rather the outcome of an act includes everything that would be the case. This includes things such as the act itself and events prior to the act. This broad notion of consequence is often accepted by consequentialist in order to develop the theory in the most ecumenical way possible.

But in another way, this notion of the outcome of an act involves certain strong commitments. In particular, it requires that there is a unique maximally specific way things would be if an agent performed an act. Very roughly, this amounts to a commitment to the principle of so-called conditional excluded middle (at least where the antecedent involves claims about what an agent does):

> $P \mathbin{\square\!\!\rightarrow} Q$ or $P \mathbin{\square\!\!\rightarrow} \neg Q$

In the standard semantics for counterfactuals, conditional excluded middle corresponds to the claim that for any P and any way things could be w, there is a unique closest possible world to w where P is true. There are ways to relax this assumption if we like, but we will adopt it for simplicity in what follows.[13]

The last feature to take note of is that this statement of consequentialism involves comparing an act's outcome with the outcome of refraining from doing the act. So standard consequentialism understands obligations in terms of the value of an outcome of an act and how it compares to the outcome of refraining from doing that act.

3.2 Generalizing

We can locate this specific way of determining whether an act is obligatory within a more general set of theoretical options about how values determine whether an act is obligatory.

---

[13]Here is a more general definition of an outcome where an outcome, o, is a proposition (a set of possible worlds):

> o is the outcome of $\mathcal{S}$'s doing x iff o is true at exactly those w such that if $\mathcal{S}$ were to do x, then w might obtain

If this more general characterization is in place, we face further questions. We must ask how the value of o is related to the value of each w. We do not explore this issue here (though some of our discussion below indirectly bears on this issue). Though I cannot pursue this here, the main claims in §A can be shown to hold in this more general setting if we make an additional assumption: if o_1, o_2, ..., o_n on are cells of a partition of set o, then it is not the case that o is strictly better than o_i for each i such that $1 \leq i \leq n$.

First, we can notice that consequentialism determines the deontic status of an act by considering the value of the possible world that would result if the act were performed. But we have already observed that there can be an outcome in which an act occurs that is not itself the outcome of the act. In Zimmerman's example, there is an outcome in which one accepts that is not the outcome of one accepting—the outcome in which one accepts the invitation and behaves well is not the outcome of accepting the invitation because one would behave poorly if one were to accept the invitation. Noticing this allows us to see that consequentialism is just one response to the question of how the deontic states of an act is related to the values of outcomes in which the act occurs. Table 1 summarizes a variety of positions one can take on this question. Standard consequentialism is a form of DEONTIC ACTUALISM. It tethers the

Table 1: How Are the Deontic Statuses of Acts Related to Outcomes in which the Act Occurs?

Views	$O(A)$ iff
DEONTIC MAXIMIN POSSIBILISM	the worst w where A obtains is better than $\mathcal{C}$
DEONTIC MAXIMAX POSSIBILISM	the best w where A obtains is the better than $\mathcal{C}$
DEONTIC ACTUALISM	the w that would obtain if A obtained is better than $\mathcal{C}$

deontic status of an act to the outcome of the act. But the table shows that there are other ways in which the value of an outcome might determine whether an act is obligatory.

We have an open parameter $\mathcal{C}$ in Table 1. This represents different views about what the relevant comparison class is for determining whether an act ought to be done. We saw earlier what standard consequentialism looks at the outcome of not doing A, but again we can imagine other answers to this question such as the ones mentioned in Table 2. BETTER THAN ALT relies on the notion of alternative which we can define

Table 2: What Comparison Class is Relevant to Determining the Deontic Status of an Act?

Views	The comparison class for A is
BETTER THAN NOT	the outcome(s) relevant to $\neg A$'s deontic status
BETTER THAN ALT	the outcome(s) relevant to every alternative to A's deontic status

as follows: B is an alternative to A for an agent exactly if A is an act available to the agent, B is act available to the agent, but the agent is not able to do $A \wedge B$. It is, of course, easy to see that neither of these tables exhausts the logical space of options.[14] Rather they simple represent a few natural ideas.

---

[14]For example, Table 2 is parasitic on what choice a theory makes in Table 1 so that there is a kind of matching evaluation between A and its comparison class. But another logically

Different choices about these issues lead to different results concerning the status of our principles. Let us look at this.

3.3 Some Relations between the Choice Points and Fault Lines

Let's begin by assuming, as the standard consequentialist does, that DEONTIC ACTU-ALISM is true and consider what if anything is interesting about the choice of between BETTER THAN NOT and BETTER THAN ALT.

As we have seen, standard consequentialism accepts DEONTIC ACTUALISM and BETTER THAN NOT. And standard consequentialism is incompatible with INHER-ITANCE, AGGLOMERATION, and NO CONFLICTS. Interestingly, however, the minor variant of standard consequentialism that accepts DEONTIC ACTUALISM but adopts BETTER THAN ALT is compatible with and indeed entails AGGLOMERATION and NO CONFLICTS. The variant of standard consequentialism that accepts BETTER THAN ALT does however share standard consequentialism's commitment to rejecting INHER-ITANCE.

Though I leave an informal proof of this to Proposition 4 in §A, we can see why we do not get a failure of these principles by returning to Zimmerman's case. There we saw that standard consequentialism entails that one is obligated to accept the invitation and behave well because this leads to the best outcome. And we saw that standard consequentialism entails that it is not the case that one is obligated to accept because the outcome of this act was worse than the outcome of not accepting. All of these claims are true according to a variant that accepts BETTER THAN ALT rather than BETTER THAN NOT as well. Since the outcome of accepting and behaving is the very best one, it is better than the outcome of every alternative. Since the outcome of accepting is worse than the outcome of not accepting, the outcome of accepting is worse than some alternative. Since both views claim that one is obligated to accept the invitation and behave well and that one is not obligated to accept the invitation, both reject Inheritance.

But standard consequentialism entails that one is obligated to decline the invitation because the outcome of declining the invitation is better than the outcome of (not not) accepting the invitation. This is verdict is what leads to the failure of AGGLOMERATION because the conjunction of declining the invitation and accepting the invitation and behaving is not obligatory. And this verdict also leads to the failure of NO CONFLICTS.

But if we adopt BETTER THAN ALT, none of these results follow: the fact that the outcome of declining is better than the outcome of accepting is not sufficient to establish that one is obligated to decline. Instead, what would need to be shown is that the outcome of declining is better than the outcome of every alternative to declining. But it is easy to see that there is an alternative to declining that has a better outcome. In particular, accepting and behaving is an alternative to decline accepting that has a better outcome. Thus, AGGLOMERATION and NO CONFLICTS are fault lines that separate DEONTIC ACTUALISM paired with BETTER THAN NOT (standard consequentialism) from DEONTIC ACTUALISM paired with BETTER THAN ALT.

---

possible approach is to use unmatched comparisons. For example, there is a theory that accepts DEONTIC MAXIMAX POSSIBILISM with a comparison class to the outcome of $\neg A$ (rather than compared to the best outcome in which $\neg A$). Though worthy of further exploration, I do not consider these unmatched approaches here.

If we now turn to Deontic Maximin Possibilism when paired with Better than Not or Better than Alt, we see that this same pattern repeats itself. As Proposition 5 in §A demonstrates, Deontic Maximin Possibilism paired with Better than Not does not validate Inheritance, Agglomeration, or No Conflicts. But, as Propoisition 7 in §A shows, Deontic Maximin Possibilism paired with Better than Alt does validate Agglomeration and No Conflicts even though it does not validate Inheritance (Proposition 6 in §A)

This repeated pattern suggests the following conjecture: any (reasonable) theory that accepts Better than Alt validates Agglomeration. To prove this conjecture, we would need a more systematic grasp of the logical space of reasonable theories. Since I cannot provide such a systematic account here, I cannot prove this conjecture. But we have encountered some circumstantial evidence for it.

Corroborating this conjecture further, Deontic Maximax Possibilism paired with Better than Alt, validates Agglomeration and No Conflicts (as is shown in Proposition 10 in §A).[15] Interestingly, however, Deontic Maximax Possibilism paired with Better than Alt standouts among the theories in that it validates Inheritance as well (as is shown in Proposition 9 in the appendix). Finally, if we turn to Deontic Maximax Possibilism paired with Better than Not, it turns out to be equivalent to Deontic Maximax Possibilism paired with Better than Alt (as is shown in Proposition 8 in the appendix)

Thus, it appears Inheritance is a fault line separating Deontic Maximax Possibilism from the other theories discussed in Table 1 and Agglomeration is a fault line separating theory that accept Better than Alt from other theories.

3.4 Further Generalizations

Though we do not have the space here to explore these issues in depth, it is worth pointing out that there are still other important answers to the two questions about how deontic statuses are determined that we are exploring. If we return to the question asked in Table 1, there are at least the further options in Table 3 to consider. Deontic

Table 3: Further Options: How Are the Deontic Statuses of Acts Related to Outcomes in which the Act Occurs?

Views	$O(A)$ iff
Deontic Maximin **F**-ism	the worst w that is **F** and where A obtains is better than $\mathcal{C}$
Deontic Maximax **F**-ism	the best w that is **F** and where A obtains is the better than $\mathcal{C}$
Deontic Averagism	some kind of average of the values of the ws that entail A is greater than $\mathcal{C}$

Maximin **F**ism and Deontic Maximax **F**-ism are like their Possibilism counterparts except that they place some further conditions on what the outcome where A obtains

---

[15]The interested reader can also check that the proof strategy for each of these results is remarkably similar: each proof exploits the fact that $\neg A \wedge B$ and $A \wedge \neg B$ are alternatives to A and B respectively to establish that $O(A \wedge B)$ given $O(A)$ and $O(B)$.

must be like. Some theories in the literature that fit this mold are so-called *securantism* (Portmore 2011 and Ross 2012) and maximalism (Portmore 2019). These views in their Maximax form are typically thought to validate Inheritance, Agglomeration, and No Conflicts. So in this respect, they pair with their Possibilism counterparts. It may also be the case that rule consequentialist views (Hooker 2001) can be thought of as forms of Deontic Maximax **F**-ism though more care is required to make this assessment.

Deontic Averagism is view that is most familiar in decision-theoretic contexts. In standard decision theory, the utility assigned to an act and whether the act ought to be done is a function of a probabilistically weighted average of the values of the various outcome in which the act it occurs. It is an interesting question for further research what properties Deontic Averagism has.

Similarly, there are are least further options in Table 4 concerning comparison to consider. Better than Threshold corresponds to simplistic forms of satisficing

Table 4: Further Options: What Comparison Class is Relevant to Determining the Deontic Status of an Act?

Views	The comparison class for A is
Better than Threshold	the outcome, τ
Better than Context	the outcome(s) relevant to the deontic status of acts supplied by context (speaker, assessor, etc.)

consequentialism (Slote 1984). And certain kinds of deontic logics that are based on preferences (Hansson 2001).

It is much harder to evaluate Better than Context without developing a much richer account of the role of context in determining comparison classes of acts. Luckily, there is already some important work in the literature about the semantics and logic of 'ought' and related notion that addresses some of these questions.[16]

This, then, gives us a number of avenues for future research by which we can assess and evaluate various fault lines and a variety of important theories that determine the deontic status of act by the value of outcomes.[17]

3.5   The Principles

Let us close this section by discussing the plausibility of the various principles that I have mentioned. In deontic logic, these principles are valid in the so-called Standard Deontic Logic. This is not to say that they are uncontroversial. Far from it. Indeed, the Standard Deontic Logic is widely rejected for a variety of reasons. I will not rehearse

---

[16]The role of context arises in discussion of cases like Zimmerman's in Jackson and Pargetter 1986. But see Cariani 2013, Cariani 2016, and Snedegar 2017 for contemporary discussion.

[17]There is a further question as to the relationship between the value of an act and the value of outcome. Most theorists accept Value Actualism according to which the value of an act is determined by the value of its outcome. But some authors (e.g., Wedgwood 2009) do not accept this view.

the challenges to Standard Deontic Logic here.[18] Instead, I will briefly illustrate the plausibility of these principles by showing how they provide tidy explanations of mundane facts about what we are obligated to do and of mundane feature of our ethical thinking. This will provide some (defeasible) evidence for the principle and show how accepting or rejecting these principle is connected to broader issues.

Begin with INHERITANCE. This principle is manifest in ordinary forms of reasoning such as concluding from the fact that you ought to drive less that fifty miles per hour on a street that ought to refrain from driving fifty-three miles per hour on that street (cf. Cariani 2013: n. 1). It is also closely related to forms of reasoning about what means we ought to take to our ends.

Next consider AGGLOMERATION. Suppose one knows that one ought to fight in the army or perform public service and suppose one also knows that one ought to not fight in the army. In this setting, it is natural to conclude that one ought to perform public service (cf. Horty 1993: 73). This inference is not licensed by INHERITANCE alone. But if one accepts AGGLOMERATION, it follows from these two claims that one ought to both fight in the army or perform public service and not fight in the army. From this and INHERITANCE, it follows that one ought to perform public service. More generally, this form of reasoning is closely related to the idea that one should consider how best to achieve one's goals taken together rather than separately.

Thus, these principles, though controversial, nicely explain simple facts about obligations and simple features of ethical thought.

4   Conclusion

Let us close by briefly mentioning how the ideas that we have explored here bear on the topic of consequentializing, the topic of whether (and how) any non-consequentialist theory can be given a consequentialist interpretation. There is no consensus about what follows from the fact that a non-consequentialist theory can be given a consequentialist representation. But Jamie Dreier provides one influential answer:

> by consequentializing a theory we can keep clearer about what the important structural differences are among competing moral theories. If I am right that the consequentialist/nonconsequentialist distinction is a shallow matter of book-keeping, then consequentializing all competitors will help shine the light on distinctions that are important, like centeredness and perhaps causal versus constitutive connections between act and consequence, by clearing away the shallow differences (Dreier 2011: 115)

According to Dreier, important differences between ethical theories are really differences in what contributes to the value of outcomes. I believe the work we have done here cast significant doubt on Dreier's idea.[19]

To see why, consider what Dreier's idea suggests about the difference between standard consequentialism and DEONTIC MAXIMAX POSSIBILISM. According to Dreier, the difference that these theories have about whether to consider the outcome of

---

[18]See McNamara 2019 and especially Hilpinen and McNamara 2014 for synoptic discussions of Standard Deontic Logic, challenges to various principles including the ones discussed here, and a broader sense of the state of play in deontic logic.

[19]It may be that Dreier's remarks are not intended to apply to the context discussed below. If so, my remarks are not a direct criticism of his view. That said, I believe that discussion here provides materials for objecting to the use of his ideas in the context he clearly intends. See also Hurley, this volume for a distinct criticism of Dreier's idea.

A or the best outcome in which *A* in determining whether an act is obligatory is not a very important one. Rather, it is better to translate the Deontic Maximax Possibilism into a standard consequentialist format and consider what theory of value it is committed to. As we have seen, such a translation would require commitment to Value Non-Splintering.

Though I have no argument for this, I put it to the reader that this does not capture what is importantly different between standard consequentialism and Deontic Maximax Possibilism. It would be uncharitable to assume that anyone believes Value Non-Splintering. And I see no reason why the matter should not be taken simply at face value: these theorists can agree about the value of outcomes but disagree about how the value of various outcomes is relevant for determining the deontic status of acts.

The lesson may generalize. What many theories disagree about is not the value of outcomes are but the way the value of outcomes is related to the deontic status of acts. These are genuine differences that, as we have seen, can correspond to important general principles.

Thus, I conclude the project of finding and assessing fault lines in ethical theory suggests that certain strands of thought in the debate about consequentializing are mistaken. And more generally, I hope that our discussion points the way to new and interesting questions that may help us to better assess long standing debates in moral theory.

A   Proofs

The proofs in this appendix are somewhat informal and rely on certain assumptions. We assume throughout that each act has a unique outcome, that the goodness ordering is a total ordering, that there is always a unique outcome that is better than all other outcomes, and that

> $\neg O_\perp$: If *A* is inconsistent, then $\neg O(A)$ is true

Relaxing these assumptions is beyond the scope of this paper. But for those who are interested, exploring these issues in a more general setting involves considering the relation between certain logical principles and certain properties of preference-like relations. There is excellent work on this topic in the deontic logic tradition.[20] But there is still room for considerable new research.

Before we turn to the proofs of Proposition 1-10, let us restate the key theses discussed in these claims (sometimes in a form more amenable for the proofs to come):

> Inheritance: If *A* entails *B*, then $O(A)$ entails $O(B)$
>
> Agglomeration: $O(A) \land O(B) \to O(A \land B)$
>
> No Conflicts: If $O(A_1)$ is true, $O(A_2)$ is true , ..., and $O(A_n)$ is true, then $\{A_1, A_2, \ldots, A_n\}$ is consistent
>
> Standard Consequentialism: $O(A)$ iff w_A is better than $w_{\neg A}$ (where w_A is the outcome of doing *A*; similarly for other acts)
>
> Deontic Actualism + Better than Alt: $O(A)$ iff w_A is better than w_B for each *B* that is an alternative to *A*

---

[20]See especially Goble 1990a, Goble 1990b, Hansson 2001.

DEONTIC MAXIMIN POSSIBILISM + BETTER THAN NOT: $O(A)$ iff w_A^- is better than $w_{\neg A}^-$ (where w_A^- is the worst outcome in which A occurs; similarly for other acts)

DEONTIC MAXIMIN POSSIBILISM + BETTER THAN ALT: $O(A)$ iff w_A^- is better than w_B^- for each B that is an alternative to A

DEONTIC MAXIMAX POSSIBILISM + BETTER THAN NOT: $O(A)$ iff w_A^+ is better than $w_{\neg A}^+$ (where w_A^+ is the best outcome in which A occurs; similarly for other acts)

DEONTIC MAXIMAX POSSIBILISM + BETTER THAN ALT: $O(A)$ iff w_A^+ is better than w_B^+ for each B that is an alternative to A

VALUE NON-SPLINTERING: if A is true at w_i and at w_j, then for any w_k where A is false, if w_i is strictly better than (/worse than/equally good as) w_k, then w_j is better than (/worse than/equally good as) w_k (where w_i, w_j, w_k are outcomes)

We now can prove each of the claims from the main text.

Proposition 1. STANDARD CONSEQUENTIALISM *and* VALUE NON-SPLINTERING *entail* INHERITANCE

Proof. Suppose A entails B and $O(A)$. Since $O(A)$, STANDARD CONSEQUENTIALISM entails w_A is better than $w_{\neg A}$. Since A entails B, B is true at w_A. By VALUE-NON-SPLINTERING, w_B is better than $w_{\neg A}$.

Suppose further that $\neg O(B)$. STANDARD CONSEQUENTIALISM entails w_B is not better than $w_{\neg B}$ (i.e., w_B is equally good or worse than $w_{\neg B}$). Since A entails B, $\neg A$ is true at $w_{\neg B}$. By VALUE NON-SPLINTERING, w_B is not better than $w_{\neg A}$ (i.e., w_B is equally good or worse than $w_{\neg A}$). Contradiction.

Thus, $O(B)$ must be true so INHERITANCE holds. $\square$

Proposition 2. STANDARD CONSEQUENTIALISM *and* VALUE NON-SPLINTERING *entail* AGGLOMERATION

Proof. Suppose $O(A)$ and $O(B)$. By STANDARD CONSEQUENTIALISM, w_A is better than $w_{\neg A}$ and w_B is better than $w_{\neg B}$, Since A and B are both true at $w_{(A \wedge B)}$, VALUE NON-SPLINTERING entails that $w_{(A \wedge B)}$ is better than $w_{\neg A}$ and $w_{(A \wedge B)}$ is better than $w_{\neg B}$.

Suppose further $\neg O(A \wedge B)$. By STANDARD CONSEQUENTIALISM, $w_{(A \wedge B)}$ is not better than $w_{\neg(A \wedge B)}$ (i.e, $w_{(A \wedge B)}$ is either equally good or worse than $w_{\neg(A \wedge B)}$). Now either (i) $\neg A$ is true at $w_{\neg(A \wedge B)}$ or (ii) $\neg B$ is true at $w_{\neg(A \wedge B)}$. Suppose (i). Then by VALUE NON-SPLINTERING, $w_{(A \wedge B)}$ is not better than $w_{\neg A}$. Contradiction. Suppose instead (ii). Then by VALUE NON-SPLINTERING $w_{(A \wedge B)}$ is not better than $w_{\neg B}$. Contradiction.

Thus, $O(A \wedge B)$ must be true so AGGLOMERATION holds. $\square$

Proposition 3. STANDARD CONSEQUENTIALISM *and* VALUE NON-SPLINTERING *entail* NO CONFLICTS

Proof. Assume $O(A_1)$, $O(A_2)$, ..., $O(A_n)$ are true but $\{A_1, A_2, \ldots, A_n\}$ is inconsistent. By Proposition 2 and $O(A_1)$, $O(A_2)$, ..., $O(A_n)$, $O(A_1 \wedge A_2 \wedge \cdots \wedge A_n)$. But given $\neg O_\perp$ and the assumption that $\{A_1, A_2, \ldots, A_n\}$ is inconsistent, $\neg O(A_1 \wedge A_2 \wedge \cdots \wedge A_n)$. Contradiction.

Thus, NO CONFLICTS holds. $\square$

Proposition 4. DEONTIC ACTUALISM + BETTER THAN ALT *entails* AGGLOMERA-TION

Proof. Suppose $O(A)$ and $O(B)$ but $\neg O(A \wedge B)$. Given DEONTIC ACTUALISM + BETTER THAN ALT, there is some act C that is an alternative to $A \wedge B$ such that $w_{(A \wedge B)}$ is not better than w_C. Since C is an alternative to $A \wedge B$, it follows either that (i) $\neg A \wedge C$ is true at w_C or that (ii) $\neg B \wedge C$ is true at w_C.

Suppose (i).Then since w_C is one where $\neg A \wedge C$ is true, $w_{(\neg A \wedge C)} = w_C$.[21] Since $O(A)$ and $\neg A \wedge C$ is an alternative to A, DEONTIC ACTUALISM + BETTER THAN ALT entail w_A is better than $w_{(\neg A \wedge C)} = w_C$. Now either $A \wedge B$ is true at w_A or it is not.

Suppose $A \wedge B$ is true at w_A. It follows that $w_A = w_{(A \wedge B)}$ so $w_{(A \wedge B)}$ is better than w_C. This contradicts our earlier claim that $w_{(A \wedge B)}$ is not better than w_C .

Suppose instead $A \wedge B$ is not true at w_A. It follows $w_{(A \wedge \neg B)} = w_A$. Since $A \wedge \neg B$ is an alternative to B and $O(B)$, DEONTIC ACTUALISM + BETTER THAN ALT entails w_B is better than $w_{(A \wedge \neg B)} = w_A$. Either $A \wedge B$ is true at w_B or it isn't. If it is, then $w_B = w_{(A \wedge B)}$ and so $w_{(A \wedge B)}$ is better than w_A which is better than w_C.This contradicts the assumption that $w_{(A \wedge B)}$ is not better than w_C. So it must be that $A \wedge B$ is false at w_B. So $w_B = w_{(\neg A \wedge B)}$. Since $\neg A \wedge B$ is an alternative to A and $O(A)$, DEONTIC ACTUALISM + BETTER THAN ALT, w_A is better than $w_{(\neg A \wedge B)} = w_B$ which contradicts our assumption that w_B is better than $w_{(A \wedge \neg B)} = w_A$. Thus (i) cannot hold.

Suppose instead, then that (ii) holds. Analogous reasoning shows that (ii) cannot be true.

Thus, Agglomeration holds. $\qquad\square$

Corollary 4.1. DEONTIC ACTUALISM + BETTER THAN ALT *entails* NO CONFLICTS

Proof. Since we have established Proposition 4, NO CONFLICTS also holds for reasons analogous to those given in Proposition 3 $\qquad\square$

Proposition 5. DEONTIC MAXIMIN POSSIBILISM + BETTER THAN NOT *does not validate* INHERITANCE, AGGLOMERATION, *or* NO CONFLICTS

A counter model. Consider a four world model where the numbers represent the value of each world:

$w_{(A \wedge B)}$: 100 $\qquad$ $w_{(A \wedge \neg B)}$: 25

$w_{(\neg A \wedge B)}$: 50 $\qquad$ $w_{(\neg A \wedge \neg B)}$: 50

Since $w^-_{(A \wedge B)}$ is better than $w^-_{\neg(A \wedge B)}$, DEONTIC MAXIMIN POSSIBILISM + BETTER THAN NOT entails $O(A \wedge B)$. But w^-_A is not better than $w^-_{\neg A}$, so DEONTIC MAXIMIN POSSIBILISM + BETTER THAN NOT entails $\neg O(A)$. Thus, INHERITANCE does not hold.

Furthermore, since $w^-_{\neg A}$ is better than w^-_A, DEONTIC MAXIMIN POSSIBILISM + BETTER THAN NOT entails $O(\neg A)$. We have already seen that $O(A \neg B)$. Since $\{A \wedge B, \neg A\}$ is inconsistent, NO CONFLICTS does not hold.

Finally, given $\neg O_\perp$, $\neg O(A \wedge B \wedge \neg A)$. Thus, AGGLOMERATION does not hold. $\quad\square$

Proposition 6. DEONTIC MAXIMIN POSSIBILISM + BETTER THAN ALT *does not validate* INHERITANCE

---

[21]This claim is justified by the assumption that if $P \boxright Q$ and $P \boxright R$, then $P \wedge Q \boxright R$ and the assumption that if P and Q are logically equivalent, then $P \boxright R$ iff $Q \boxright R$. These same assumptions are used in a number of other places in this proof.

A counter model. Consider the same model as the one given for Proposition 5.

$w^-_{(A \wedge B)}$ is better than all worlds, so $w^-_{(A \wedge B)}$ is better than w^-_C for any C that is an alternative to $A \wedge B$. So by Deontic Maximin Possibilism + Better than Alt, $O(A \wedge B)$.

But once again w^-_A is not better than $w^-_{\neg A}$ and $\neg A$ is an alternative to A, so Deontic Maximin Possibilism + Better than Alt entails $\neg O(A)$. Thus, Inheritance does not hold. $\square$

Proposition 7. Deontic Maximin Possibilism + Better than Alt *entails* Agglomeration

Proof. We begin with a lemma.

Lemma 7.1. Deontic Maximin Possibilism + Better than Alt *entails that if* $O(A)$ *and* $O(B)$*, then either* A *entails* B *or* B *entails* A*.*

Proof of lemma. Suppose A does not entail B and B does not entail A. Then there is a $w^-_{(\neg A \wedge B)}$ and there is a $w^-_{(A \wedge \neg B)}$.

Suppose $O(A)$ and $O(B)$. By Deontic Maximin Possibilism + Better than Alt and the fact that $\neg A \wedge B$ is an alternative to A, w^-_A is better than $w^-_{(\neg A \wedge B)}$. And similarly, w^-_B is better than $w^-_{(A \wedge \neg B)}$.

But if w^-_A is better than $w^-_{(\neg A \wedge B)}$, then $w^-_{(A \wedge \neg B)}$ is better than $w^-_{(\neg A \wedge B)}$. And similarly, $w^-_{(\neg A \wedge B)}$ is better than $w^-_{(A \wedge \neg B)}$. This is a contradiction and so proves the Lemma. $\square$

We now return to the main proof. Suppose $O(A)$ and $O(B)$, but $\neg O(A \wedge B)$. Given Deontic Maximin Possibilism + Better than Alt, there is an alternative C to $A \wedge B$ such that $w^-_{(A \wedge B)}$ is not better than w^-_C.

By Lemma A either A entails B or B entails A. Suppose A entails B so A and $A \wedge B$ are equivalents and so every alternative to A is an alternative to $A \wedge B$ and vice-versa and $w^-_A = w^-_{(A \wedge B)}$. Since $O(A)$, Deontic Maximin Possibilism + Better than Alt entails that w^-_A is better than w^-_C for any C that is an alternative to A. Thus, $w^-_{(A \wedge B)}$ is better than w^-_C for any C that is an alternative to $A \wedge B$. This is a contradiction.

So suppose B entails A. Analogous reasoning shows we reach a contradiction.

Thus, Agglomeration holds. $\square$

Corollary 7.1. Deontic Maximin Possibilism + Better than Alt *entails* No Conflicts

Proof. Since we have established Proposition 7, No Conflicts also holds for reasons analogous to those given in Proposition 3 $\square$

Proposition 8. Deontic Maximax Possibilism + Better than Alt *and* Deontic Maximin Possibilism + Better than Not *give equivalent verdicts about what is obligatory*

Proof. This claim trivially follows from Lemma 8.1 and Lemma 8.2 below. $\square$

Lemma 8.1. *If* $O(A)$ *is true according to* Deontic Maximax Possibilism + Better than Alt*, then* $O(A)$ *is true according* Deontic Maximin Possibilism + Better than Not

Proof. Suppose $O(A)$ is true according to Deontic Maximax Possibilism + Better than Alt. Thus, w_A^+ is better than w_C^+ for every alternative C to A. $\neg A$ is an alternative to A. Thus w_A^+ is better than $w_{\neg A}^+$ is true. So according to Deontic Maximax Possibilism + Better than Not, $O(A)$ is true. $\square$

Lemma 8.2. *If $O(A)$ is true according to* Deontic Maximax Possibilism + Better than Not*, then $O(A)$ is true according to* Deontic Maximin Possibilism + Better than Alt

Proof. Suppose $O(A)$ is true according to Deontic Maximax Possibilism + Better than Not. Thus w_A^+ is better than $w_{\neg A}^+$. Now for any w_C^+ such that C is an alternative to A, $\neg A$ is true at w_C^+. So $w_{\neg A}^+$ is at least as good as w_C^+. Thus for any alternative C to A, w_A^+ is better than w_C^+. So according to Deontic Maximax Possibilism + Better than Alt, $O(A)$ is true. $\square$

Proposition 9. Deontic Maximax Possibilism + Better than Alt *entails* Inheritance

Proof. Suppose A entails B and $O(A)$. By Deontic Maximax Possibilism + Better than Alt, w_A^+ is better than w_C^+ for any alternative C to A. Since A entails B, B is true at w_A^+. Therefore w_B^+ is at least as good as w_A^+. Next if C' is an alternative to B, C' is an alternative to A since A entails B. Thus, w_A^+ is better than $w_{C'}^+$ for any alternative C' to B. So since w_B^+ is at least as good as w_A^+, w_B^+ is better than $w_{C'}^+$ for any alternative C' to B. So by Deontic Maximax Possibilism + Better than Alt, $O(B)$. Therefore, Inheritance holds. $\square$

Proposition 10. Deontic Maximax Possibilism + Better than Alt *entails* Agglomeration

Proof. We begin with a lemma.

Lemma 10.1. Deontic Maximax Possibilism + Better than Alt *entails that if $O(A)$ and $O(B)$, then $w_A^+ = w_{(A \wedge B)}^+ = w_B^+$*

Proof of lemma. Assume $O(A)$ and $O(B)$. Given Deontic Maximax Possibilism + Better than Alt and $O(A)$, w_A^+ is better than w_C^+ for any alternative C to A. Either $A \wedge B$ is true at w_A^+ or it isn't. If $A \wedge B$ is true at w_A^+, then $w_A^+ = w_{(A \wedge B)}^+$.

Suppose instead $A \wedge B$ is false at w_A^+. Thus, $A \wedge \neg B$ is true at w_A^+ and $w_A^+ = w_{(A \wedge \neg B)}^+$. Since $A \wedge \neg B$ is an alternative to B and $O(B)$, Deontic Maximax Possibilism + Better than Alt entail that w_B^+ is better than $w_{(A \wedge \neg B)}^+ = w_A^+$. Now either A is true at w_B^+ or it isn't.

Suppose A is true at w_B^+. It follows that since w_B^+ is better than $w_{(A \wedge \neg B)}^+$, $w_{(A \wedge \neg B)}^+ \neq w_A^+$. This contradicts our previous claim that $w_A^+ = w_{(A \wedge \neg B)}^+$.

Suppose then A is false and so $w_{(\neg A \wedge B)}^+ = w_B^+$. Since $\neg A \wedge B$ is an alternative to A and since $w_{(\neg A \wedge B)}^+ = w_B^+$ is better than w_A^+. So Deontic Maximax Possibilism + Better than Alt entails $\neg O(A)$ which contradicts our assumption that $O(A)$. So $A \wedge B$ must be true at w_A^+. So $w_{(A \wedge B)}^+ = w_A^+$.

Next, given Deontic Maximax Possibilism + Better than Alt and $O(B)$, w_B^+ is better than w_C^+ for any alternative C to B. Either $A \wedge B$ is true at w_B^+ or it isn't. By analogous reasoning, we establish $A \wedge B$ is true at w_B^+ and therefore $w_{(A \wedge B)}^+ = w_B^+$.

Thus, $w_A^+ = w_{(A \wedge B)}^+ = w_B^+$. $\square$

We return now to the main proof. Suppose $O(A)$ and $O(B)$, but $\neg O(A \wedge B)$. By DEONTIC MAXIMAX POSSIBILISM + BETTER THAN ALT, there is C that is an alternative to $A \wedge B$ such that $w^+_{(A \wedge B)}$ is not better than w^+_C. Since C is an alternative to $A \wedge B$, either (i) $\neg A \wedge C$ is true at w^+_C or (ii) $\neg B \wedge C$ is true at w^+_C.

Suppose (i). So $w^+_C = w^+_{(\neg A \wedge C)}$. Given $O(A)$ and the fact that $\neg A \wedge C$ is an alternative to A, DEONTIC MAXIMAX POSSIBILISM + BETTER THAN ALT entails that w^+_A is better than $w^+_{(\neg A \wedge C)} = w^+_C$. But by Lemma 10.1, $w^+_A = w^+_{(A \wedge B)}$ so $w^+_{(A \wedge B)}$ is better than $w^+_{(\neg A \wedge C)} = w^+_C$. This contradicts our earlier claim that $w^+_{(A \wedge B)}$ is not better than w^+_C. So (i) must be false.

Suppose (ii). Analogous reasoning establishes that (ii) must be false.

Thus, AGGLOMERATION holds. $\square$

Corollary 10.1. DEONTIC MAXIMAX POSSIBILISM + BETTER THAN ALT *entails* NO CONFLICTS

Proof. Since we have established Proposition 10, NO CONFLICTS also holds for reasons analogous to those given in Proposition 3 $\square$

References

Brown, Campbell (2011). "Consequentialize This". In: *Ethics* 121 (4), pp. 749–771.

— (2018). "Maximalism and the Structure of Acts". In: *Noûs* 52 (4), pp. 752–771.

Cariani, Fabrizio (2013). "'Ought' and Resolution Semantics". In: *Noûs* 47 (3), pp. 534–558.

— (2016). "Consequence and Contrast in Deontic Semantics". In: *Journal of Philosophy* 113 (8), pp. 396–416.

Darwall, Stephen (1986). "Agent-Centered Restrictions from the Inside Out". In: *Philosophical Studies* 50 (3), pp. 291–319.

Dietrich, Franz and Christian List (2017). "What Matters and How It Matters". In: *Philosophical Review* 126 (4), pp. 421–479.

Dreier, James (1993). "Structures of Normative Theories". In: *The Monist* 76 (1), pp. 22–40.

— (2011). "In Defense of Consequentializing". In: *Oxford Studies in Normative Ethics*. Ed. by Mark Timmons. Vol. 1. Oxford: Oxford University Press, pp. 97–119.

Foot, Phillipa (1967 [1979]). "The Problem of Abortion and the Doctrine of Double Effect". In: *Virtues and Vices*. Berkeley: University of California Press, pp. 19–32.

Goble, Lou (1990a). "A Logic of *good*, *should*, and *would* Part I". In: *Journal of Philosophical Logic* 19, pp. 169–199.

— (1990b). "A Logic of *good*, *should*, and *would* Part II". In: *Journal of Philosophical Logic* 19, pp. 253–276.

Hansson, Sven Ove (2001). *The Structure of Values and Norms*. New York: Cambridge University Press.

Hilpinen, Risto and Paul McNamara (2014). "Deontic Logic". In: *Handbook of Deontic Logic and Normative Systems*. Ed. by Dov Gabbay et al. Vol. 1. Milton Keynes: College Publications, pp. 3–136.

Hooker, Brad (2001). *Ideal Code, Real World*. Oxford: Oxford University Press.

Horty, John (1993). "Deontic Logic as Founded on Nonmonotonic Logic". In: *Annals of Mathematics and Artificial Intelligence* 9 (1-2), pp. 69–91.

Jackson, Frank and Robert Pargetter (1986). "Oughts, Actions, and Actualism". In: *The Philosophical Review* 95 (2), pp. 233–255.

Louise, Jennie (2004). "Relativity of Value and the Consequentialist Umbrella". In: *The Philosophical Quarterly* 54 (217), pp. 518–536.

McNamara, Paul (2019). *Deontic Logic*. In: *Stanford Encyclopedia of Philosophy*. Ed. by Edward Zalta. Summer 2019. URL: https://plato.stanford.edu/archives/sum2019/entries/logic-deontic/.

Nair, Shyam (2014). "A Fault Line in Ethical Theory". In: *Philosophical Perspectives* 28 (1), pp. 173–200.

Oddie, Graham and Peter Milne (1991). "Act and Value". In: *Theoria* 57 (1-2), pp. 42–76.

Portmore, Douglas (2007). "Consequentializing Moral Theories". In: *Pacific Philosophical Quarterly* 88 (1), pp. 39–73.

— (2009). "Consequentializing". In: *Philosophy Compass* 4 (2), pp. 329–347.

— (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.

— (2019). *Opting for the Best*. New York: Oxford University Press.

Ross, Jacob (2012). "Actualism, Possibilism, and Beyond". In: *Oxford Studies in Normative Ethics*. Ed. by Mark Timmons. Vol. 2. Oxford: Oxford University Press, pp. 243–282.

Scheffler, Samuel (2003 [1982]). *The Rejection of Consequentialism*. Revised. Oxford: Oxford University Press.

Schroeder, Mark (2007). "Teleology, Agent-Relative Value, and 'Good'". In: *Ethics* 117 (2), pp. 265–295.

Slote, Michael (1984). "Satisficing Consequentialism". In: *Proceedings of the Aristotelian Society* 58, pp. 139–163.

Snedegar, Justin (2017). *Contastive Reasons*. Oxford: Oxford University Press.

Wedgwood, Ralph (2009). "Intrinsic Values and Reasons for Action". In: *Philosophical Issues* 19 (1), pp. 321–342.

Williams, Bernard (1973). "A Critique of Utilitarianism". In: *Utilitarianism: For and Against*. Ed. by J.J.C. Smart and Bernard Williams. Cambridge: Cambridge University Press, pp. 75–150.

Zimmerman, Michael (2006). "The Relevant Risks to Wrongdoing". In: *The Good, The Right, Life And Death: Essays in Honor of Fred Feldman*. Ed. by Kris McDaniel et al. Aldershot: Ashgate, pp. 151–172.