

Consequentialism and Our Best Selves

Preprint— Please Cite the Published Version (*Philosophical Studies*
<https://doi.org/10.1007/s11098-022-01879-2>)

Abstract: I develop and defend a maximizing theory of moral motivation: I claim that consequentialists should recommend only those desires, emotions, and dispositions that will make the outcome best. I advance a conservative account of the motives that are possible for us; I say that a motive is an alternative if and only if it is in our *psychological control*. The resulting theory is less demanding than its competitors. It also permits us to maintain many of the motivations that we value most, including our love for those most important to us. I conclude that we are closer to meeting morality's demands on our character than has been appreciated.

Keywords: consequentialism; motives; demandingness; alternatives; partiality

Morality's demands may be divided in two. First, we may be required to perform certain actions. These demands can be expressed as claims about what we should *do*. Second, we may be required to possess certain mental states, dispositions, or states of character. These demands can be expressed as claims about how we should *be*. I am a consequentialist about what we should do; I believe that:

We should perform some act if and only if (and because) it would bring about the best outcome.

What does consequentialism tell us about how we should be?

Perhaps the most fundamental question of how we should be is the question of how we should be motivated. I understand our motivations broadly: they are the parts of our psychology (including our desires, emotions, and dispositions) that can move us to act. Many assume that:

If we should do only what will make things best, then we should also be motivated only to do what will make things best.

However, as Williams (1973) notes, a person who desires only to maximize the impartial good—a kind of ideal consequentialist—is a worrying figure. He develops friendships just to improve the

outcome, rescues his wife to maximize value, and is uninterested in developing his talents and projects, as there is always more good to be done.¹

So understood, consequentialism presents an unattainable and unattractive ideal. If we should be motivated only to bring about the best outcome, then we cannot *also* be motivated by anything else. But few of us could replace all of our motivations with an unyielding desire to maximize value. And even if we could make these changes, they would be profoundly costly; we would need to abandon the emotions and desires that help make us who we are. Indeed, any morality that asks us to reject our most valued motivations, including our love for our family and friends, is deeply implausible.² When we think of how we might become better persons, we do not think that we should strive never to be motivated by love again.

I aim to develop an alternative account of how we should be motivated. I hope for a view that is consistent with consequentialism but respects our limitations and our humanity. Such a view will not ask us for what we cannot give. It will not make extreme demands of our psychology. And it will, in general, allow us to maintain our most valued motivations, including our love for those most important to us.

In what follows, I first argue that no traditional theory of moral motivation can achieve these goals. I consider right-making and Kantian views; I claim that when combined with consequentialism, these theories require motives that are impossible and undesirable. After, I examine the most common alternative accounts of moral motivation offered by consequentialists such as Moore (1993) and Railton (1984). Though these views make more plausible recommendations and demand less of us, they are untenable too: they ignore the independent difference our motives can make to the outcome.

I argue instead that we should evaluate our motives directly and endorse a *maximizing theory of moral motivation*. According to this theory, we ought to possess some motive if and only if it would bring about the best outcome. Such views have been suggested by Parfit (1984), Feldman (1993), and Pettit and Smith (2000) but have not been sufficiently developed; in particular, these philosophers do not say what our alternative motives *are* and so provide little guidance about how much morality may demand of our character. I advance a conservative view; I claim that a motive is

¹ See Williams (1973), (1981), Railton (1984), and Wolf (1982) respectively. These concerns have more general application; see Baron (1984) and Piper (1987) for discussion. I think Mason (1999) is correct, however, that “the mud seems to have stuck to consequentialism in particular.”

² Jackson (1991) writes only that it is “chilling” to imagine that morality might require that we jettison our fundamental commitments to our friends and family.

an alternative for us if and only if it is within our *psychological control*. So understood, consequentialism's evaluation of our character is attractive: it does not require the impossible or make otherwise excessive demands. And it may allow us to preserve the motives that we value most.

§1: The Standard Accounts of Moral Motivation

I believe that no standard theory of moral motivation can be successfully combined with consequentialism. This may seem unsurprising: these theories are usually advanced by those who endorse some form of moderate deontology. But the problem that arises indicates something important about the relation between consequentialist theories of right action and the account of moral motivation that I believe we should endorse.

I begin with *the right-making view*; it declares:³

We are motivated as we should be if and only if (and because) we are motivated by the features that make actions right.

The right-making view reflects the idea that the best moral agents are those who do the right thing, for the right reason. So, if Lucy should walk Linus to school because she has promised to do so, then she should be moved by her desire to keep her promise—rather than, say, a desire to avoid being punished by her parents. Or if it would be right for Sally to give some of her Halloween candy to Charlie Brown because it would make him happy, then it is the fact that her act would make Charlie Brown happy that should move her—not a desire for recognition or reward.

The right-making view is simple and attractive. But it justifies the link between motivation and right action that I had hoped to sever: if the right-making view is true, then the consequentialist should be motivated only to maximize the good. As a result, morality will ask us to be motivated in ways that are unattractive and impossible.

The right-making view is also inconsistent with the spirit of consequentialism.⁴ Imagine that because Lucy is motivated to maximize the impartial good, she systematically acts so as to make things worse: she takes on grand responsibilities that she mismanages, interferes in others' lives and

³ My formulation of the right-making view is grounded in remarks by Arpaly (2003), Markovits (2010), and Stratton-Lake (2011). That said, these philosophers present the right-making view as a position about virtuous or morally worthy motivation, not what motivations we *should* have. I assume, however, that these philosophers would agree that, in general, we should be motivated in ways that are virtuous—indeed Stratton-Lake (2011: 373) suggests that the motivations the right-making view recommends are those that are morally “required” (though compare Markovits (2010: 235)).

⁴ See Parfit (1984: 24).

activities, and neglects her friends and family.⁵ If we believe that we should act to make the outcome *best*, we cannot insist that we should also be motivated in ways that will make the outcome *worse*.

The problem is deeper than it first appears. The moral evaluation of our actions cannot be grounded in features that do not matter.⁶ In particular, it cannot be true that:

(i) our acts should make the outcome as good as possible

even though

(ii) it does not matter if the outcome is as good as possible.⁷

The consequentialist cannot, of course, abandon (i) and so will reject (ii) instead. But if it is morally important that the outcome be as good as possible, then we should reject motives that make the outcome worse. This conclusion could be avoided only if there is something else that might be *more important* than maximizing the good. But the existence of some further, overriding moral goal would undermine the truth of consequentialism.

So the right-making view cannot be combined with consequentialism. But these same problems affect the other standard account of moral motivation, *moralism*. According to the moralist (or, as some call it, *Kantian*) position:

We are motivated as we should be if and only if (and because) we are motivated by rightness *per se*.

Again, this position is grounded in an attractive idea: an agent who cares more about doing what is right is better than one who cares less. Thus, the best agent would be motivated always and exclusively to do what is right. Lucy should be motivated to walk Linus to school not because she wishes to keep her promise, but because she wishes to fulfill her obligations. Similarly, Sally should

⁵ Stratton-Lake (2011: 373) presents a similar argument to show that the right-making view and consequentialism are inconsistent but concludes that this is simply so much the worse for consequentialism. See also Markovits (2010: 234-236).

⁶ On this point see especially Portmore (2019: 1-14). (I disagree with Portmore, however, that what “ultimately matters” is what we should non-instrumentally care about.) I should stress that there is a difference between what ultimately matters *morally* and what ultimately matters *epistemically*, *aesthetically*, and so on. I therefore do not assume that all things must be evaluated in terms of what matters morally or that the consequentialist must e.g. reject the legitimacy of epistemic normativity.

⁷ Pettit and Smith (2000) argue similarly.

not be concerned with Charlie Brown's happiness *per se*, but with meeting morality's requirements—with doing what she should.

However, like the right-making view, moralism's demands are unreasonable: we cannot—and should not—replace all of our motives with the desire to do what is right. Further, in cases like Lucy's, moralism will recommend motives that make things worse. It is therefore also inconsistent with the fundamental goal of consequentialist morality.

§2: A Revision to the Right-Making View

The consequentialist must reject the most familiar theories of moral motivation. Could a revised version of these views do better? Chappell (2021) presents an attractive and novel revision of the right-making view; he claims it may be combined successfully with consequentialism. I take his position to be the best candidate.

Chappell argues that the right-making features of moral theories have been misunderstood. Most assume that we can simply 'read off' the right-making features of a theory by examining its account of right action. But this is a mistake:

Our theories may advert to highly abstract properties in specifying their criteria for right action: that which fills in the blank in statements of the form, "An act is right iff ." But we need not take those canonical criteria to themselves be the theory's fundamental moral grounds. *Instead, I propose, we should interpret them as summarizing the full range of moral grounds posited by the theory.* Highly abstract summary criteria are compatible with appropriately concrete and personal ground-level concerns. (2021: 427; emphasis mine)

Consequentialists claim that acts are right *if and only if* they maximize the good. But it does not follow that *maximizing the good* is the fundamental right-making feature of actions. Rather, Chappell claims, *maximizing the good* is a derivative, general property; the actual right-making features of acts are specific—they are the 'ground-level' features of acts that give us reasons to perform them. How might these features be identified?

The general answer, I think, involves looking to the particular morally significant components (or normative reasons) that explain how the general criteria for rightness come to be satisfied in any given case. For act consequentialism, the general criterion is producing the best outcome, and the components that contribute to the satisfaction of this criterion are

particular token values (specified in the theory's axiology). Consequentialists should presumably hold that agents have pro tanto reason to ϕ just when ϕ -ing promotes some token value... So, if we believe that virtuous agents would be moved by concern for particular individuals (and not just the general good), consequentialists can accommodate this by adopting a token-pluralistic axiology that specifies each individual's welfare as a separate basic good. (2021: 438)

So, imagine that Sally makes things best by giving Charlie Brown her candy. According to Chappell, Sally's action is right not because it maximizes the good, but because, in performing it, she promotes a particular token value—namely, Charlie Brown's happiness.

This reply allows that consequentialists should be motivated by a plurality of concerns. We are not required to possess an overriding drive to maximize the good; we can instead be moved by a desire to benefit those we love. As a result, we can understand the ideal consequentialist not as someone dedicated to maximizing the good, but as someone dedicated to the happiness of individuals (as well as any other token values that we identify). Further, in requiring only that we desire to benefit individuals, morality asks for less; indeed, many of us already possess these kinds of desires. This is a powerful advantage.

Yet I believe that consequentialists should reject Chappell's theory. Either the more specific properties Chappell appeals to make actions right or the more general property of *maximizing the good* makes actions right—but not both. (This would double our reasons to act rightly.) If we must choose, consequentialists should choose general right-making features, not specific ones; it is these general features that matter.

Consider an analogy from the metaphysics of causation. Imagine that a bird has been trained to peck at all and only red patches. One day, the bird sees a red patch, and it pecks. Suppose further that this patch is a particular shade of red, say, crimson. We have then two possible causes: the bird pecked because the patch is red or the bird pecked because the patch is crimson.⁸

We should not say both are causes; this would be double counting. Intuitively, it is the patch being red that matters: if the patch were not red, the bird would not have pecked. But if the patch were not crimson, but rather scarlet, the bird still would have pecked.

⁸ This example is due to Yablo (1992). See also his (1997). I am thankful to Bradford Skow for suggesting an appeal to proportionality to settle these conflicts between moral explanations.

This is Yablo's (1992) test of *proportionality*: we should appeal to the cause whose generality matches the generality of the effect. This is the property that makes the difference. Counterfactual reasoning reveals that *being crimson* is too specific. But *being red* is just right.⁹

Though the relation between rightness and right-making features is non-causal, the same reasoning applies. When explanations compete, we should choose the explanation that is proportional to what is being explained. Counterfactual reasoning reveals that *maximizing the good* is the proportional explanation. If an action maximized the good but did not benefit Charlie Brown it would, of course, still be right—but the reverse is not true. *Benefiting Charlie Brown* is, for the consequentialist, like *being crimson*: it has no intrinsic significance.

Still, for the sake of argument, imagine that we admit the specific right-making features Chappell identifies and jettison the general right-making features I prefer. Ask instead: *why* should we possess the motives that Chappell recommends? Suppose that, in being motivated only by the specific right-making properties to which Chappell appeals, I would not in fact promote the token values he identifies. If we care most about promoting these values, we must say that I should be motivated differently. In this way, Chappell's view is like the standard right-making and moralist positions: all may recommend motivations that make things worse; all are inconsistent with the spirit of consequentialism.

§3: Against Fittingness

Chappell may reply that I have misunderstood his theory:

An agent may ask, "What motivational profile would it be best or most desirable, from a moral perspective, for me to have?"... This is to ask about the morally recommended or fortunate motivational profile. By contrast, the question I'm interested in is importantly different—something more along the lines of, "What motivational profile is most morally fitting or apt, reflecting an orientation toward the good, and is on this basis perhaps worthy of praise or high esteem?" (2021: 428).¹⁰

I have assumed that the fundamental question about the moral status of a motive is whether we should possess that motive. Chappell rejects this assumption; he suggests that, even if we are

⁹ The test of proportionality is in fact somewhat more complicated; see Woodward (2010) for discussion. However, for our purposes, this simple form of the test will suffice.

¹⁰ See also his (2021: 431, fn. 11).

consequentialists, we must allow that there are two distinct moral features of a motivation: whether we *should* possess it and whether it would be *fitting* to possess it.¹¹ He may therefore agree that his view does not answer the question of how we should be motivated. But it does answer an important question about moral motivation that I have ignored: which motivations are fitting or appropriate?

I am skeptical of the moral significance of fittingness. Some part of our psychology, like an emotion or disposition, may be fitting regardless of its moral status.¹² Indeed, as Fritz (2021) argues, there are many possible events that are both *bad* and *uncertain* in a way that makes it fitting to be anxious about them. But morality in no way *prefers* that we be anxious about every likely bad state of affairs; such facts about fittingness generate no moral reasons.

As a result, in proclaiming that some emotion or desire would be fitting to have, we have not answered any question about the *moral status* of that emotion or desire. Similarly then, in answering the question of which motivations are fitting or apt, we have not answered any question about the moral status of these motivations.

Suppose, however, that there is a particular concept of *moral* fittingness and that some motivations are morally fitting. I ask: what is the normative role of this concept?

Chappell (2012) argues that fittingness is *sui generis*.¹³ To say that some attitude is morally fitting does not imply that we should have that attitude, or that it would be good to possess that attitude. Thus facts about whether a motive is morally fitting seem irrelevant to the value of the outcome. So how, then, does the concept of a fitting motive earn its keep? It seems to be an idle wheel in an otherwise austere consequentialist moral ontology.¹⁴

Here Chappell has a powerful reply.¹⁵ Imagine that a demon will destroy the world unless I possess some motive that would otherwise be wrong, such as an intrinsic desire for others to suffer. I *should* have this desire, given the alternatives. But this motivation is surely flawed. In desiring the suffering of others, we *get it wrong*. This motivation is not fitting and thus does not reflect an orientation towards the good. In this way the concept of fittingness is needed to describe the moral

¹¹ See his (2012), (2021).

¹² D'Arms and Jacobson (2000) thus warn of a 'moralistic fallacy' when theorizing about fittingness.

¹³ Note that if fittingness did have some input to what we should do but was not relevant to the value of things, then this would be inconsistent with the truth of consequentialism.

¹⁴ Chappell (2012) argues also that fittingness might play a fundamental *definitional* or *structural* role in a consequentialist ontology: we might analyze goodness in terms of fittingness (see also McHugh and Way (2016)). I reject fitting attitude analyses of goodness. I am a Moorean about the structure of moral philosophy; I have attempted to defend my views in Tucker (2020). However, I think this issue is ultimately tangential to the question I am interested in here: namely, whether *in addition to questions about the deontic status of a motive* there is an independent and morally important question about the *fittingness of a motive*.

¹⁵ Chappell (2021: 428).

complexity of this situation.¹⁶ Without it, there seems to be nothing unfortunate here: in wishing for the suffering of others we are simply being the kind of person that morality recommends.

I agree that the situation described is morally complex. But I am unconvinced that this complexity requires a concept of moral fittingness. Consider an analogy. Imagine that a demon will destroy the world unless I kill an innocent person. I *should* kill this innocent person, given the alternatives. Yet the act of killing a person seems to be flawed also; we do not wish to say that this is a case of happily satisfying morality's demands. So, by parallel reasoning, this act of killing must be unfitting.

Yet consequentialists rarely give such explanations; they present a more familiar apology.¹⁷ Killing an innocent person will bring about many bad things—the destruction of a life and all of its connections to the good. This is a terrible thing; we understandably recoil from it. However the badness of this killing is outweighed by the *far worse things* that must be prevented. This is why the consequentialist believes we should comply with the demon. And it is also why the situation is unfortunate. For the consequentialist, what is important is the value of the outcome; there is nothing *intrinsically good* about fulfilling our obligations. Thus this situation should not be understood as one in which we happily comply with morality's demands. Rather it is one where, no matter what we do, something bad will happen. This is, from the consequentialist's perspective, a deeply unfortunate situation.

I believe this explanation generalizes. If I desire that others suffer, this will bring about many bad things: it will warp or destroy what is most important to me, including my relationships with those I love. It will also prevent many good things by eliminating my capacity to enjoy the happiness of others. As a result, this situation is unfortunate: it is a choice between two evils. There are no residual facts that fittingness must explain.

Chappell might reply that my defense fails: the intuition that I would be *getting it wrong* in desiring the suffering of others remains, even if nothing bad would come about from my possessing this desire. I am unsure that I share this intuition. But I think it can be explained regardless—and I think it can be explained without the concept of fittingness. I say that we get it wrong in possessing the desire that others suffer because desire *presents its objects as good*.¹⁸ Thus, an intrinsic desire for the suffering of others presents the suffering of others as intrinsically good. But suffering is perhaps the

¹⁶ See Chappell (2012).

¹⁷ See e.g. Smart (1973: 71).

¹⁸ For defense of this thesis see e.g. Oddie (2017).

most profound intrinsic evil. We get it wrong when we have this desire, because we think something deeply false about morality.¹⁹ This thought is not unfitting, it is merely *false*.

Still, I expect that some readers may be unsympathetic to my rejection of fittingness. The issues are complex and important; I cannot advance my conclusions further here. So suppose instead that I am mistaken and that there is an independent, important question about the moral fittingness of our motives. I would allow then that Chappell's theory succeeds on its own terms. But I would insist that it does not answer the most pressing question for the consequentialist: which motives should we have? Morality's *demands* are deontic; fittingness is not about what morality demands or recommends, as Chappell notes.²⁰ So, if we wish to say with confidence that morality does not demand that we abandon the motives we value most, this is the question we must answer.

§4: The Broad Moorean View

I turn now to views of moral motivation that have often been advanced by consequentialists; I shall claim that, surprisingly, we should reject these views too.

In the *Principia Ethica*, G.E. Moore advanced a hierarchical normative structure: goodness is fundamental; right acts are those that maximize the good; virtues are states of character that produce right actions:²¹

[A] virtue may be defined as an habitual disposition to perform certain actions, which generally produce the best possible results. Nor is there any doubt as to the kind of actions which it is 'virtuous' habitually to perform. They are, in general, those which are duties... (1993: 221).

Moore's view does not apply to all motives but only to habitual dispositions. And Moore speaks not of the motives we *should* have but of the motives that are *virtuous* to have. However, if we allow that

¹⁹ In speaking of what we 'think' I do not mean what we 'believe'. Thoughts may take the form of 'striking' or 'seemings' with propositional content; they can therefore be true or false. See Howard (2018) and Rosen (2015). We might say that when the thoughts implicit in some attitude are false, then this attitude is unfitting in some merely descriptive sense. I have no objection to this proposal; I expand upon it in fn. 35. But I note that this descriptive notion of fittingness is radically different than the *sui generis* notion postulated which is, of course, normative—this is why the question "which motives is it fitting to have" is supposed to be a question about *moral* motivation. See Howard (2018: 12, en. 30).

²⁰ Chappell (2021: 428).

²¹ See Moore (1993: 192, 196-197, 219-220).

other motivations can be virtuous—and that we *should be* virtuously motivated—we can generate a broader and more powerful view.²² According to this extended view:

We are motivated as we should be if and only if (and because) we are motivated in ways that will lead to our acting so as to maximize the good.

The broad Moorean position fits neatly with the pragmatic character of consequentialism. Unlike the standard right-making and moralist views, the Moorean position ignores the *content* of our motives and focuses only on their effects. If it would make things best for Lucy to walk Linus home, then it does not matter if she is moved by love, moral concern, or self-interest—she is rightly motivated as long as her motives lead her to do what will bring about the best.

The Moorean view is nearly identical with the various accounts of moral motivation advanced in the works of “sophisticated consequentialists” such as Railton (1984).²³ These philosophers argue that consequentialism permits any set of motives, as long as those motives are consistent with living a life in which we perform “the most beneficial overall sequence of acts” (1984: 160).²⁴

Such accounts generate a more agreeable conception of the ideal consequentialist. There is no reason to believe that being motivated always to do what is best will, in fact, lead to doing so. Rather, if beings like us possessed such motives, we would likely act in ways that would be worse than if we were motivated in more humble, human ways. As a result, the Moorean view also seems less demanding than its competitors.

The Moorean view is therefore attractive. Still, I think we should reject it. The view is incomplete: it evaluates our motives only in terms of how they affect our actions. But our motives can make a difference to the outcome, independent of our actions. Imagine a wife discovers that her husband’s acts of care are grounded only in obligation, or a patient discovers that his doctor is driven only by a desire for medical knowledge. In such cases, our motives can make a difference to how things will be. Part of this difference may manifest in our actions. But even if these motivations do not change how we act, they can change the outcome by changing how we feel—and how we relate to each other.

²² See fn. 3.

²³ See also Norcross (1997), Mason (1998), (1999), and Conce (2001).

²⁴ In his (1988) Railton suggests a different view, *valoric utilitarianism*, which evaluates all things directly, in relation to the outcome. However, this version of Railton’s view does not assign deontic evaluations; it is closer to a kind of global, scalar consequentialism.

The problem is compounded by the fact that the Moorean view must be understood to range over *intentional actions*; unintentional actions cannot be “duties” or even alternatives.²⁵ Thus the Moorean principle should be understood to claim that we should possess whatever motives would lead to our performing some *intentional act* that maximizes the good. But our motives can, of course, affect our *involuntary* or *unintentional* actions. And these involuntary actions can affect the outcome (imagine how a spy’s nervous, unintentional movements may affect an interrogation).

I therefore argue that if:

- (i) our motives can make an independent contribution to the outcome

and:

- (ii) if consequentialism is true, then what ultimately matters is the value of the outcome

then it must be true that:

- (iii) if consequentialism is true, we cannot evaluate our motives only in terms of how they affect our actions.

If we did, then the independent contribution that our motives make to the outcome would be lost. And it seems unjustifiable to exclude this difference simply because it cannot be traced back to our voluntary actions.

A fantastical case may establish the same point.²⁶ Suppose that, because of a demon’s curse, we cannot act but can be motivated in certain ways (i.e. there are various desires, emotions, or dispositions available to us). Suppose further that through unknown connections (perhaps the powers of the demon or merely unusual natural laws), our motives consistently affect what will be. In such cases, the broad Moorean view declares that all of our motives and states of character are equally good, since our motivations cannot lead us to act rightly. Yet this is surely the wrong verdict.²⁷

²⁵ See Portmore (2019: 40).

²⁶ I am thankful to Mikhail Valdman for making this point to me.

²⁷ We might worry that this is a misunderstanding of the broad Moorean view; it should be formulated (as perhaps Moore intended) in terms of what is *generally* the case—where ‘generally’ is a modal restriction that eliminates unusual worlds, like the demon world. According to this modal version of Moore’s view, we are rightly motivated when our motivations are such that, in normal worlds, these motivations would lead us to act to produce the best result. This will do little, however, to help us avoid the objection. Suppose the demon will inflict terrible pain and suffering upon all of us if we are motivated in ways that, in normal worlds, lead us to act in ways that make things best. Surely, then, *if we are in*

§5: Conflicts Between Motives and Actions

If we agree that:

- (i) our fundamental goal is that the outcome be as good as possible
- (ii) acts and motives are simply different ways of affecting the outcome
- (iii) we should perform those acts that make the outcome best

Then we should accept:

The Maximizing Theory of Moral Motivation: We are motivated as we should be if and only if (and because) we possess some motive that will make the outcome best.²⁸

This view engenders none of the problems mentioned so far. It does not demand the impossible, since it selects the motives we should possess from those that are available to us; it fits neatly with the spirit of consequentialism; and it avoids the incompleteness of the broad Moorean view.

Moreover, I shall argue that it makes attractive, plausible claims about morality's demands on our character.

Still, the prospects of the maximizing theory may seem bleak. This is because, as Adams (1976) notes, it allows for morality to require actions and motives that are jointly impossible.²⁹

How might these conflicts arise? Consider Parfit (1984: 32)'s example of Clare. Clare deeply loves her child. She must choose between benefiting her child or giving a larger benefit to a stranger. According to consequentialism, it is likely wrong for Clare to benefit her own child.³⁰ But she does so anyway. Yet Clare acted wrongly only because of her strong love for her child and (let us stipulate) this is the best motive available to her. Further, it would not be possible for her to be motivated by her love and also to benefit the stranger.

Thus the best actions and the best motives are incompatible. And if we ought to perform the best actions and possess the best motives, then it seems that the maximizing theory violates the principle that 'ought' implies 'can'.

the demon world, we should not possess these otherwise valuable motives: they would do nothing but make things as bad as possible, while helping no one.

²⁸ My argument is therefore similar to a special case of the powerful arguments advanced in Pettit and Smith (2000) and Smith (2008). See also Greaves (2020).

²⁹ See also Streumer (2003); for reply see Brown (2005).

³⁰ Of course, as Parfit (1984: 32) mentions, we might endorse some complicated form of consequentialism that attributes intrinsic value to love *from one's own parents*. Such views might have different implications in the case given, but we could simply modify the significance of the benefit so as to overcome this.

I take this problem seriously. I suggest two possible replies.

§5.1: Ought to be and Ought to do

According to the first reply, though it may be true that:

- (i) Clare ought to benefit the stranger and
- (ii) Clare ought to be motivated by her deep love for her child

it does not follow that:

- (iii) Clare ought to both benefit the stranger and be motivated by her deep love for her child.

The ‘ought’ in claims such as ‘you ought to possess some motive’ expresses the *ought to be*. However, the ‘ought’ in claims such as ‘you ought to perform some action’ expresses the *ought to do*. There is, I believe, no further ‘ought’ that takes these as inputs. Thus, the claim ‘Clare ought to *both* benefit the stranger and be motivated by love for her child’ cannot be true: ‘ought’ must either express the *ought to be* or the *ought to do*, but neither can take both actions and motives as arguments.³¹ We therefore cannot generate any violation of the principle that ‘ought’ implies ‘can’.

As a result, morality will require a kind of *evaluative ambivalence*.³² Situations like Clare’s will pull us in two directions—we will describe such cases as a kind of moral failing and a kind of moral success. Clare acted wrongly; she did not make things best. But this claim will now be tempered: she acted wrongly only because of her motivations, which were the best available to her.

³¹ We might object that the *ought to do* is, in fact, reducible to the *ought to be*—roughly, we ought to do something when it ought to be that we do it. I think this reply cannot succeed. I take it that claims about what ought to be divide into two types. Some things ought to be *relative to an agent*—as when we say that Lucy ought to be kinder. Some things ought to be *simpliciter*—as when we say that it ought to be that no one dies of hunger. (I expand on this distinction in §6.) But the *ought to do* cannot be reduced to either the relativized or non-relativized *ought to be*. Begin with the relativized notion. Consider the claim that we ought to do something when *we ought to be* such that we do it. This claim is false; it may be best for us to be the kinds of persons that sometimes act wrongly. So turn to the non-relativized notion: we claim that we ought to perform some action when it ought to be *simpliciter* that we perform that action. But this claim is false, too. To say that something ought to be *simpliciter* is to say, roughly, that it would occur in a world that is morally ideal. But we may be obligated to act in certain ways (such as in cases of self-defense) only because others *will in fact* act wrongly; in such cases we ought to do something (such as kill an attacker) though, in a morally ideal world, we would not act in this way. Thus, from the fact that some act is obligatory it does not follow that it ought to be. See Feldman (1986: 179-196).

³² See Driver (2012: 149). In fact, Driver’s position is richer than the one I suggest here; she imagines cases in which our actions and motives are wrong, but our motives are nonetheless *good*, allowing for further complexity. I think this is possible also, but I will not investigate such ambivalence here.

I do not find these verdicts worrying or inconsistent. Quite the opposite: I think these ambivalent judgements are plausible because they reflect the moral complexity of Clare's situation. Yet I admit that this kind of ambivalence can complicate morality's guidance.

The concern must be stated carefully. We cannot merely object that when the best motives and best actions conflict, we will not know how we ought to act or how we ought to be motivated. These questions will always be answered for us by the theories proposed. Rather, the concern is that there is some further question:

When the best actions and the best motives conflict, is it *more morally important* that I perform the best act or possess the best motive?

It is unclear that such questions can be answered. This may seem damaging—indeed it may seem to defeat the maximizing theory. But if so, it defeats nearly every alternative.

Consider some Kantian theory according to which morality recommends that I be motivated always to do what is right *per se* but says also that I ought to perform those actions that comply with the categorical imperative. Further suppose that, due to my unusual psychology, though I can possess these motives and can perform these actions, it is impossible for me to do *both*. In such cases, the morally best motives will conflict with my moral obligations. We may ask: is it morally more important to perform the actions that I ought to perform or to possess the motives that I ought to have? It is unclear how these questions could be answered.

Such problems are not rare and will not be limited to conflicts between motivations and actions: these puzzles will arise whenever morality recommends two things—whether they be actions, intentions, motives, beliefs, or lives—that are incompatible. Further, such conflicts need not be deontic. We can ask: is it more important to live a good life or to act in ways that are praiseworthy? Is it more important to be virtuous or to act rightly? In such cases, morality's guidance may come to an end: we may be told how to act, how to be motivated, how to live—but not how to weigh such things against each other.

I do not think this is objectionable. Still, if we are unsatisfied with this *tu quoque*—or if my intuitions about how morality must guide us are not shared—then a second defense of the maximizing theory is available.

§5.2: World Consequentialism

My second reply appeals to the powerful *world consequentialism* of Feldman (1986), (1993). Feldman's view harmonizes morality's demands by eliminating any fundamental moral distinction between motives and actions. Rather, it understands *all* of morality's requirements as part of our overall obligation to make the world as good as it can be. While Linus might be obligated to help Lucy with her homework and obligated to be motivated by a desire for her wellbeing, these are fundamentally the *same* kind of obligation. Specifically, they are obligations *to make propositions true*. Linus is obligated to make it true that he helps Lucy with her homework and to make it true that he is motivated by his desire for her wellbeing. And Linus will have these obligations because these propositions are true at the best possible world accessible to him.

This reply avoids any violation (apparent or otherwise) of the principle that 'ought' implies 'can': it tells us to perform those actions and possess those motives (understood as certain classes of propositions) that we do and have in the best world accessible to us. The accessibility relation takes the place of the 'can' of personal possibility—and whatever restrictions we build into the latter, we may build into the former, enabling an easy simulation of the maximizing view I have proposed.

If Clare cannot see to it that she both benefits the stranger and that she is motivated by her love for her child, then there will be no world accessible to her in which these propositions are true. Whether she should then benefit the stranger or be motivated by her love will depend upon the values of the worlds where these propositions hold. If the best accessible world is one where she benefits the stranger, then this is what she should see to. If the best accessible world is instead one where she is motivated by her love, then this is what she should see to.

Thus, we have two attractive accounts of how the maximizing view might be combined with the principle that 'ought' implies 'can'; either will suffice for my purposes. I conclude that the maximizing view is not refuted by concerns about conflicts between the best actions and the best motives.

§6: The Maximizing View and Morality's Demands

I return to the nature of the maximizing theory. To say that a motive is best is to say that there is no alternative that is better. But when is one motive an alternative to another? In answering this question, we set the boundaries of how much the maximizing view can demand. The more we are capable of, the more morality can ask for.

Though I cannot develop a complete account here, I advance a conservative theory. Such a theory will diverge significantly from Adams (1976), who suggests that an individual's alternative motives are those motives that it is possible for humans *in general* to possess. Act consequentialism respects our limited agency; it does not say that some action is an alternative for me because humans can generally perform it. A consequentialist theory of motives should respect our limitations and reject the idea that a motivation is available for me simply because others could possess it.

According to the conservative view I prefer, whether a motive is within our power is entirely a matter of which psychological states are within our power. In particular, a motive is an alternative for me when it is in my *psychological control*.

Psychological control divides into two categories. A mental state is in my *direct* psychological control when I can immediately bring it about. I assume that states such as willing, intending, imagining, and reflecting are often within our direct control. A mental state is in my *indirect* control when I can bring it about by means of mental states that are in my direct control.

Our motives are, I believe, in our indirect control. Suppose that though Sally does not wish to give any of her candy to Charlie Brown, she could be motivated to do so by imagining how he must feel and reflecting on their relationship. If so, being motivated to give away her candy is an alternative for Sally.

In this way Sally's power over her motives is like our power over our non-basic actions. We may have control over the direction of an automobile in virtue of having control over how we move the steering wheel; we may have control over the structure of an essay by having control over the keys we type. Similarly, Sally may have power over her motives in virtue of her power over her will, her imagination, and her reflective capacities.

This view about our alternatives is restrictive. But it still allows for a rich conception of morality's demands on our character. Consider my attempt to be motivated to care for others as I care for myself. I may begin by trying to care for others only a little, by vividly imagining their circumstances and reflecting on what others have done for me. If this is all I can do at a time, then I satisfy morality's demands by doing so. But once I care for others a little, new alternatives appear. My ability to understand others deepens, allowing me to be motivated to care more deeply for them. The ideal shifts forward. At each stage, as I grow, morality asks for more from me.

This picture implies that morality's demands may be *continual*. If we can be better, then we must be; there is no point at which we might claim that we have become 'good enough.' I do not

find this objectionable. For beings like us, the project of becoming a good person is unending, regardless of the progress we have made.

Further, the demands of a theory may be continual without being excessive. Consider a theory of right action that, for any time, assigns a deontic status to every action available to us at that time. This theory will make continual demands; in every circumstance there will be *something* that we should do (even if it is a disjunction). But this does not mean that this theory is too demanding. (Imagine a theory that says that we should always do whatever would be easiest for us.)

Of course, like all plausible moral theories, complying with the maximizing view *will* sometimes be costly. But its demands will be reasonable. The maximizing theory will not require an unrelenting focus on grand self-improvement: we are limited creatures and this would only make the outcome worse. We will do better by instead focusing on more moderate changes that are clearly within our power.

Further, when we consider the outcome of some motive, we must consider how it affects *us*. Even if there is some available motive that might benefit others, this does not mean that we must adopt this motive, especially if it would be difficult or painful for us to do so. Indeed, we may sometimes simply be too exhausted, stubborn, or selfish to possess motives that would otherwise be better. At these times, morality asks for nothing more. In this way, consequentialism's demands are sensitive to what is possible for us.

This may seem to be a double-edged sword, however—in refusing to demand more of us, the maximizing theory may seem to condone viciousness. Imagine that Lucy is cruel: she lives a selfish, narrow life and hates to see others thrive. But imagine also that she is incapable of being different. Must we be silent about her flaws?

Such agents are rare—at the very least, most of us can be *less* cruel. But necessarily flawed agents are possible. Though these persons meet morality's demands, they contribute to making outcomes worse. This is a kind of moral flaw, even if it is one that they cannot change. Is such a flaw a sufficient basis for moral criticism or blame?

Consequentialists—and utilitarians in particular—have often claimed that we should criticize some person when it would make things best, regardless of whether this person has complied with morality's demands.³³ There is no further question, they insist, about whether the objects of our

³³ See e.g. Smart (1973: 54): “Whose was *the* responsibility? The act-utilitarian will quite consistently reply that the notion of *the* responsibility is a piece of metaphysical nonsense and should be replaced by ‘Whom would it be useful to blame?’” For a more recent contribution to this general tradition see Morris (2017). A variant of this proposal speaks not of the value of blaming some individual but of the value of our general *practices* of praise and blame; see Mason (2020).

reproach are legitimate targets. This position need not be supported only by an impassionate utilitarian calculus; it can be grounded in the idea that persons cannot deserve to suffer. But for a person to be blameworthy, they would need to deserve to be blamed and thus feel guilt—and to feel guilt is to suffer.³⁴

Regardless of the rationale, the classical consequentialist position insists that no one is *worthy* of blame. But blaming will nonetheless be obligatory when it brings about the best result. If this is correct, morality may allow us to criticize the motives of necessarily flawed agents like Lucy. The objection dissolves.

But suppose instead we find the traditional consequentialist reply outmoded and insist that, even if we *should* blame someone when it would make things best, there is a further question about the legitimacy of this blame. We claim it is a necessary (but not sufficient) condition for a person to be a legitimate target of moral criticism that they have somehow failed to do what morality demands of them.³⁵ I accept then that Lucy is not to blame for her motives—*she* is as she ought to be. But we may say that it ought to be *that* she is different.

I distinguish between two types of the *ought to be*.³⁶ Some things ought to be *relative to an agent*—as when we tell Lucy that she ought to be less cruel. I take this kind of *ought* to be a *directive ought*; it tells Lucy how to be.³⁷ On the other hand, some things ought to be *simpliciter*—as when we say that it ought to be that no one dies of hunger. I take this to be a kind of *non-directive* (or *evaluative*) *ought*; it does not instruct anyone.

When a person has become the best that they can be, we cannot truly say that they ought to be different in the directive sense of ‘ought’; morality respects our limitations and so can never tell us to be someone that we cannot be. Nonetheless, in a non-directive sense, it ought to be that

³⁴ Parfit (2011: 163-172) denies that we can deserve to suffer. Portmore (2019: 54-63) draws a close connection between a person being blameworthy and it being fitting for them to experience guilt; he argues further that fitting guilt involves deserved suffering.

³⁵ Can I allow, however, that there is such a thing as *legitimate* moral criticism? I have expressed skepticism about fittingness as a *sui generis* moral concept. I am also skeptical of similar notions, such as desert. It may therefore seem infelicitous to speak of *legitimate* moral criticism: surely legitimate moral criticism is simply criticism that is *fitting* or *deserved*. But though I have objected to fittingness as a *sui generis* moral notion, I do not object to a concept of fittingness that is simply descriptive. According to the *alethic* view of fittingness, some moral criticism such as blame, regret, or condemnation will be fitting (or as I say, legitimate) if and only if the thoughts it implicates (such as *that person did something wrong*) are true (see Rosen (2015)). This kind of fittingness is simply a matter of representing things as they are; I do not take it to generate any moral reasons. When I speak of criticism that is legitimate, I mean to speak of criticism that is fitting only in this minimal, descriptive way. I do not think that in employing this concept, I must abandon any of the claims I have made thus far.

³⁶ I reference this distinction also in fn. 31.

³⁷ See Portmore (2019: 24).

necessarily flawed agents are different. Such persons make things worse; their character is instrumentally bad. A necessarily flawed agent therefore represents a kind of moral fault in the world.

What does this imply for reactive attitudes, such as blame? Though Lucy's cruelty may cause things that are bad, this does not make her a legitimate target of blame; she has complied with morality's directives. But because, in a non-directive sense, it ought to be that she is different, we might cast blame upon a situation that she is a part of, including the events that caused her to be unable to become better.

This is clearer in the case of action. Consider those who are compelled to perform actions that would otherwise be wrong—whether they be kleptomaniacs or those driven to directly harm others. If 'ought' implies 'can', then 'cannot' implies 'not ought.' Thus, the actions of such individuals cannot be wrong; they cannot act otherwise. Criticism therefore seems inappropriate; as Widerker (2000) argues, such agents may ask 'what should I have done instead?'³⁸ But we may focus our blame on the circumstances that caused them to be this way.

The same is true of motivation. If we criticize agents like Lucy for their motivations, then they might ask 'how should I be instead?' If we cannot answer this question, then we cannot focus our criticism on them. But, again, we may place blame upon the circumstances that help explain their psychological limitations.

We may be tempted to object that it is incoherent to speak of blaming a circumstance. But I think this objection is misguided. It is widely recognized that there are at least two kinds of blame: causal blame ("I blame the boiler for the explosion") and interpersonal blame ("I blame Lucy for the explosion"). Chislenko (2021) argues persuasively that these two kinds of blame are more similar than has been assumed. The fundamental difference between causal and interpersonal blame lies in our *focus*: in both cases we make a judgement about the causal origin of something bad. But in interpersonal cases we focus on the person *as the cause* of the bad. When our focus shifts from a person to a circumstance *involving* that person, our blame begins to resemble merely causal blame. I say only that, when a person cannot be better, it is illegitimate to focus our blame on them. Instead, we must shift our focus from the person to an object or situation.

There is surely more to say about the nature of moral criticism; I do not pretend that my remarks here settle the issue. But ultimately I believe that the problem is not unique to the

³⁸ See also Portmore (2019: 66-69) for an insightful discussion of this point.

maximizing theory. Rather, it is a general issue: what will we say about those who cannot help but make things worse? This is a difficult question, and we may differ over how best to answer it. But most answers that take our limitations seriously will, I believe, be compatible with the maximizing theory.³⁹

§7: Conclusion

I have aimed to provide a consequentialist account of moral motivation that complies with the principle that ‘ought’ implies ‘can’, makes reasonable demands, and does not suggest that we abandon our most valued motivations. When combined with the conservative view of alternatives that I have suggested, the maximizing theory satisfies these constraints.

Recall the objection suggested by critics such as Williams (1973). Consequentialism recommends a worrying ideal: someone who develops relationships just to promote the good; acts to protect others only because of the value of the outcome; and is uninterested in the tasks, projects, and commitments that structure a human life. Is such an ideal reasonable, according to the maximizing theory?

No. If we consider the consequentialist’s view of what the ideal *act* should be, we find that this concept is unclear; the deontic status of an act is necessarily tied to its alternatives and thus to the agent. There is therefore no universally ideal action but only the best action we can perform. Indeed, all right acts are, in a sense, *ideal*: they represent the best that we can do.

The same is true of our motives. The maximizing theory asks for the most we can give; it identifies the ways we should be with the ways that are ideal. We are not required to be the type of person that the critics of consequentialism describe; we do not need to replace our love for our family and friends with an overwhelming desire to maximize the good. These motives are not alternatives for us. And even if we could possess such motivations, they would often damage our close relationships and lead to our taking on nearly impossible tasks that most of us would fail at, making things worse.

³⁹ Some may reject the analogy I have drawn; they may say that criticism of actions differs fundamentally from criticism of motives. I disagree—as I have attempted to show throughout this paper, actions and motives are simply different ways of affecting the outcome and thus, for the consequentialist, their moral status should be determined in fundamentally the same way. Because I endorse this kind of deontic equivalence, I believe also that the justification of moral criticism of actions and motives should be addressed in fundamentally the same way.

Instead, according to the maximizing theory, *there is no universally ideal agent*. Rather, there is simply the best moral agent that *we* can be—the best version of ourselves. This is what morality asks of us. And it is this vision of our best self that is the correct ideal for the consequentialist.⁴⁰

⁴⁰ I am extremely grateful to Fred Feldman, Bradford Skow, Peter Graham, Robert Gruber, Andrew Moon, Mikhail Valdman, Anthony Ellis, James Fritz, Eugene Mills, Catherine Sutton, Kevin Brosnan, Sarah Colquhoun, Donald Smith, Cooper Ackerly, Stephen Ingram, Lisa Tucker, and two very helpful anonymous referees from *Philosophical Studies*.

Bibliography

- Adams, Robert. (1976) "Motive Utilitarianism." *The Journal of Philosophy*, 73,14: 467-481.
- Arpaly, Nomy (2003) *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford University Press.
- Baron, Marcia. (1984) "The Alleged Moral Repugnance of Acting from Duty." *The Journal of Philosophy*, 81,4: 197-220.
- Brown, Campbell. (2005) "Blameless Wrongdoing and Agglomeration: A Response to Streumer." *Utilitas*, 17, 2: 222-225.
- Chappell, Richard Yetter. (2012) "Fittingness: The Sole Normative Primitive." *Philosophical Quarterly*, 62, 249: 684-704.
- . (2021) "The Right Wrong-Makers." *Philosophy and Phenomenological Research*, 130, 2: 426-440.
- Chislenko, Eugene. (2021) "Causal Blame." *American Philosophical Quarterly*, 58, 4: 347-358.
- Conee, Earl. (2001) "Friendship and Consequentialism." *Australasian Journal of Philosophy*, 79, 2: 161-179.
- D'arms, Justin and Daniel, Jacobson. (2000) "The Moralistic Fallacy: On the 'Appropriateness' of Emotions." *Philosophy and Phenomenological Research*, LXI, 1: 65-90.
- Driver, Julia. (2012) *Consequentialism*. Routledge: New York.
- Feldman, Fred. (1986) *Doing the Best We Can*. D. Reidel Publishing Company: Dordrecht, Holland.
- . (1993) "On the Consistency of Act- and Motive-Utilitarianism: A Reply to Robert Adams." *Philosophical Studies* XXX.
- Fritz, James. (2021) "Fitting Anxiety and Prudent Anxiety." *Synthese*, 199: 8555-8578.
- Greaves, Hilary. (2020) "Global Consequentialism." In *The Oxford Handbook of Consequentialism*. Ed. Portmore, Douglas. New York: Oxford University Press: 423-440.
- Howard, Christopher. (2018) "Fittingness." *Philosophy Compass*. 13, e12542.
- Jackson, Frank. (1991) "Decision-theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics*, 101, 3: 461-482.
- Markovits, Julia. (2010) "Acting for the Right Reasons." *Philosophical Review*, 119, 201–242.
- Mason, Elinor. (1998) "Can an Indirect Consequentialist Be a Real Friend?" *Ethics*, 108: 386-393.

- . (1999) “Do Consequentialists Have One Thought Too Many?” *Ethical Theory and Moral Practice*, 2: 243-261.
- . (2020) “Consequentialism, Blame, and Moral Responsibility.” In *The Oxford Handbook of Consequentialism*. Ed. Portmore, Douglas. New York: Oxford University Press: 162-178.
- McHugh, Conor and Way, Jonathan. (2016) “Fittingness First.” *Ethics*, 126: 575-606.
- Morris, Rick. (2017) “Praise, Blame, and Demandingness.” *Philosophical Studies*, 174: 1857-1869.
- Moore, G.E. (1993) [1903] *Principia Ethica*. Ed. Baldwin, Thomas. Cambridge University Press: New York.
- Norcross, Alastair. (1997) “Consequentialism and Commitment.” *Pacific Philosophical Quarterly*.
- Oddie, Graham. (2017) “Desire and the Good: In Search of the Right Fit.” In *The Nature of Desire*. Eds. Deonna, Julien A. and Lauria, Federico. Oxford University Press: New York: 30-56.
- Parfit, Derek. (1984) *Reasons and Persons*. Oxford University Press: Oxford.
- . (2011) *On What Matters: Volume I*. Oxford University Press: Oxford.
- Paul, Sarah K. (2021) *Philosophy of Action: A Contemporary Introduction*. Routledge: New York.
- Pettit, Philip and Smith, Michael. (2000) “Global Consequentialism.” In *Morality, Rules, and Consequences*. Eds. Hooker, Brad, Mason, Elinor, and Miller, Dale. Edinburgh: Edinburgh University Press: 121-133.
- Piper, Adrian M.S. (1987) “Moral Theory and Moral Alienation.” *The Journal of Philosophy*, 84, 2: 102-118.
- Portmore, Douglas W. (2019) *Opting for the Best*. Oxford University Press: New York.
- Railton, Peter. (1984) “Alienation, Consequentialism, and the Demands of Morality.” *Philosophy and Public Affairs*, 13, 2: 134-171.
- . (1988) “How Thinking about Character and Utilitarianism Might Lead to Rethinking the Character of Utilitarianism.” *Midwest Studies in Philosophy*, 13, 1: 398-416.
- Rosen, Gideon. (2015) “The Alethic Conception of Moral Responsibility.” In *The Nature of Moral Responsibility: New Essays*. Eds. Clarke, Randolph, McKenna, Michael and Smith, Angela M. Oxford University Press: New York: 66-88.
- Smart, J.J.C. (1973) “An Outline of a System of Utilitarian Ethics.” In *Utilitarianism: For & Against*. Eds. Smart, J.J.C. and Williams, Bernard. Cambridge University Press: New York: 3-74.

- Smith, Michael. (2008) "Consequentialism and the Nearest and Dearest Objection." In *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*. Ed. Ravenscroft, Ian. Oxford University Press: New York: 237-265.
- Stark, Cynthia A. (1997) "Decision Procedures, Standards of Rightness and Impartiality." *Noûs* 31, 4: 478-495.
- Stratton-Lake, Philip. (2011) "Recalcitrant Pluralism." *Ratio* 24: 364–383.
- Streumer, Bart. (2003) "Can Consequentialism Cover Everything?" *Utilitas* 15: 237-247.
- Tucker, M. (2020). "Moore, Brentano, and Scanlon: A defense of indefinability. *Philosophical Studies*, 177(8), 2261–2276."
- Yablo, Stephen. (1992) "Mental Causation." *Philosophical Review*, 101, 2: 245-280.
- . (1997) "Wide Causation." *Philosophical Perspective*, 11: 251-281.
- Widerker, David. (2000) "Frankfurt's Attack on the Principle of Alternative Possibilities: A Further Look." *Philosophical Perspectives*, 14: 181-201.
- Williams, Bernard. (1973) "A Critique of Utilitarianism." In *Utilitarianism: For & Against*. Eds. Smart, J.J.C. and Williams, Bernard. Cambridge University Press: New York: 77-150.
- . (1981) "Persons, Character and Morality." In *Moral Luck*, Ed Rachels, James. Cambridge University Press: New York: 1-19.
- Wolf, Susan. (1982) "Moral Saints." *The Journal of Philosophy*, 79, 9: 419-439.
- Woodward, James. (2010) "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy*, 25: 287-318.