

The Alienation Objection to Consequentialism

Calvin Baker (Princeton University) & Barry Maguire (Stanford University)¹

Forthcoming in the *Oxford Handbook of Consequentialism*.

Penultimate Draft. Comments still welcome! Kindly cite / quote from published version.

Abstract: An ethical theory is alienating if accepting the theory inhibits the agent from fitting participation in some normative ideal, such as some ideal of integrity, friendship, or community. Many normative ideals involve non-consequentialist behavior of some form or another. If such ideals are normatively authoritative, they constitute counterexamples to consequentialism unless their authority can be explained or explained away. We address a range of attempts to avoid such counterexamples and argue that consequentialism cannot by itself account for the normative authority of all plausible such ideals. At best, consequentialism can find a more modest place in an ethical theory that includes non-consequentialist principles with their own normative authority.

Keywords: Alienation, Commitment, Motives, Normative Ideals, Authority

I. Alienation

The state of alienation is rather complex.² An agential *subject* is alienated from some *object* – perhaps a friend, a stranger, a group, or their own self³ – when they are inhibited from fitting participation in some pertinent normative ideal involving the object. The normative ideal will often involve a relationship of some kind, for instance, being a friend, a teacher, or a fellow citizen; there are also ideals concerning the relations between one’s own ethical principles, commitments, affects, and motives. You might be alienated from your friend by your acceptance of overly formal norms concerning friendly intercourse; or from your community by the prevalence of an ideology of individual achievement; or from yourself by your failure to be motivated by your own averred ethical ideals.

Not everything plausibly thought of as alienation involves this structure (cf. Schacht, 1970 and Leopold, 2018). Another important theme in historical

¹ Calvin Baker calvinbaker1@gmail.com, Barry Maguire, barry.maguire@stanford.edu. Many thanks to Tristram McPherson, Austen McDougal, Milan Mossé, Oded Na’aman, Doug Portmore, and Jack Woods for excellent comments on earlier drafts.

² For historical discussions of alienation, see Schacht, 1970 and Kain, 1982. For recent discussions see Wood, 1999; Brudney, 2014; Jaeggi, 2014; and Leopold, 2018.

³ The 1844 Marx discussed alienation from other people, but also from one’s labor, the product of one’s labor, and one’s species being. Hegel, Marx, and Fromm talked about alienation from the natural world. One might also wonder about cases of being alienated from activities, institutions, cultures, etc. We will be principally concerned with alienation from oneself or others, more or less directly.

discussions concerns the relinquishing of control: whether over political authority (e.g. in Hobbes or Rousseau) or civil society (e.g. in Hegel), or the patterns and outputs of economic and social production (e.g. in Marx), or in the aesthetic quality of our lives (e.g. in Schiller and Marcuse). However, we do think there is a core idea, strongly associated with the notion of alienation, involving the inhibition of participation in some normative ideal for a kind of relationship. This will be our focus.

We will restrict our attention to normative ideals that involve some sort of *harmony, connection, or closeness* between persons or groups. Such relationships involve standards of fittingness bearing on one's participation (cf. Wood, 1999 and Leopold, 2018). In this sense, the restricted concept of a normative ideal that concerns us involves its own distinctive normative standards. But these standards are not, as such, *authoritatively* normative, in the sense of necessarily bearing on what one ought to do, all things considered.⁴ This is clear from consideration of harmful ideals: consider the just alienated from the racist society; or the socialist alienated in the libertarian society. In such cases, there may be no authoritative reason for the alienated agent to fit into the relevant normative ideal; rather there is authoritative reason for the prevailing norms to change.⁵

It does not quite follow that alienation is, as such, undesirable. For one thing, undesirability is an axiological property and non-authoritativeness is a deontic property. It is so far an open question whether there are normative ideals the realization of which would be good but the standards of which are not authoritative, for instance if their normativity is defeated. This is not too far-fetched: realizing the second-best normative ideals might bring about a great deal of value, but nevertheless it may be that the availability of a better ideal renders the better standards authoritatively normative instead.

The more pertinent question is the following: Are there any normative ideals that are authoritative, the authority of which does not depend upon consequentialist assessment? To sharpen the question a little, say that for an ethical principle to be ethically *fundamental* is for it to be authoritative and not ethically explained by anything else (cf. Maguire, 2015). For an ethical principle to be *supreme* is for it to be the only ethically fundamental principle. If there is at least one normative ideal that is authoritative, the authority of which does not depend upon whether its realization meets the test of consequentialist assessment, then Supreme Consequentialism – the thesis that the consequentialist standard of assessment is the only fundamental ethical principle – is false. It seems extremely plausible that there is a range of normative ideals – concerning varieties of friendship, community, and integrity – that are authoritative in this way. Hence, it is a cost to a consequentialist ethical theory if it is unable to instruct agents to properly accord with such ideals.

⁴ On authoritative and non-authoritative normativity, see Maguire & Woods (fc.) and references therein.

⁵ Note that it is also an open question whether a given normative ideal is *operative* in a domain. We assume that an ideal is operative in a context only if a set of ideal-sustaining norms are accepted in the context. You might be alienated from a friend *by* the suboptimal norms of friendship that prevail in your social context, or by something else in spite of the availability of excellent friendship norms. There might be multiple norms of friendship available in a large social context, such that one would count as alienated relative to one but not another.

How much of a cost? We lack space to address this methodological question. We will just say this. It is plausible that some authoritative normative ideals constitute inputs to ethical theorizing that any satisfactory ethical theory needs to accommodate. Supreme Consequentialism does not have this default epistemic status. It purports to *be* the final ethical theory – the output of this process of ethical theorizing. Consequently, if Supreme Consequentialism conflicts with such normative ideals, it is Supreme Consequentialism that should be revised, rather than the normative ideals.

Our question, then, is whether this or that ‘consequentialist’ ethical theory yields the result that an agent ought to properly accord with this or that normative ideal. We will consider four classes of ethical theories: Hybrid Theories, Relative Value Theories, Global Consequentialism, and Levelled Consequentialism.⁶ We will argue that there are some plausibly authoritative normative ideals that cannot be accommodated by the consequentialist principle in any of them. It follows that any version of Supreme Consequentialism is alienating relative to these ideals.

II. The Alienation Challenge Facing Consequentialist Theories

By *consequentialism* we mean an ethical principle according to which pertinent deontic facts about something are fully explained by facts about the value promoted by that thing and its alternatives, or by some set of suitably related alternatives. A version of consequentialism is *direct* if deontic facts about alternatives are explained by the values promoted by those alternatives, and *indirect* if explained by the values promoted by some suitably related alternatives.⁷

Let Harmonious Act Consequentialism be the conjunction of direct Act Consequentialism, according to which one ought to perform an action if and only if and because it promotes more value than any alternative, and the principle that one ought to be motivated, in taking an option, by whatever explains why one ought to do so (cf. Arpaly, 2011 and Way, 2017).

The simplest alienation objection runs as follows. Suppose that taking some option that benefits Joe is the alternative that promotes the most value. You take the option, and you are motivated by the fact that doing so promotes more value than any alternative. Suppose that Joe is a close friend of yours. According to a plausible normative ideal of friendship, it is fitting to be motivated to help friends by the fact that one’s action will help one’s friend, not by the fact that it will promote the

⁶ We title an ethical *theory* – as opposed to an ethical principle that is a constituent in an ethical theory – “consequentialist” only if it maintains that some consequentialist principle is supreme.

⁷ As Julia Driver points out (this volume), ‘indirect consequentialism’ is used in two different ways in the literature. One is to denote an indirect deontic standard, as described in the main text. But it is sometimes used to distinguish consequentialisms employing a sufficiently robust distinction between the deontic standard itself and the correct decision-making procedure. We’ll stick with ‘indirect’ for the former, and use ‘self-effacing’ (allowing that this comes in degrees) for the latter.

most value overall.⁸ The objection goes that one's acceptance⁹ of Harmonious Act Consequentialism, insofar as it ensures that one lacks the former motive, alienates one from one's friend (cf. Williams, 1973 and Markovits, 2010).

This version of the objection moves too quickly. To see this, consider the difference between Harmonious Act Consequentialism and Harmonious Eight-Ball, which maintains that one ought to do whatever the eight-ball randomly decides, and that one ought to be motivated to do it because the eight-ball said so. In the case of Harmonious Act Consequentialism, but not Harmonious Eight-Ball, facts about the benefits to one's friend are part of the *full* explanation of why one ought to benefit Joe, if not part of the *immediate* explanation for why one ought to do so. So, facts about the benefit to Joe will be part of one's motivation in performing the action, alongside facts about the various other costs and benefits of doing so.

This may not yet be fully satisfying. Every fact about the anticipatable effects of your alternatives on anything evaluatively significant until the end of time will show up in your motivational structure, *qua* part of the explanation of why you ought to perform this action. But another refinement is available to the Harmonious Act Consequentialist: they can acknowledge a distinctive level of normative explanation, between facts about values and facts about oughts, consisting of facts about reasons; and they might distinguish between reasons and their background conditions (cf. Schroeder, 2007 and Maguire, 2016). The facts about actions promoting states of affairs might motivate an agent on the condition that those states of affairs are valuable, and in proportion, other things being equal, to their value. The distinction between one's motivating reason and its background condition helps to ensure that some facts about the individual being helped feature in the content of specific motives; the proportionality condition elevates the significance of the impacts on the beneficiary above other parts of the full explanation. One might then add a pragmatic story, again along the lines sketched in Schroeder, 2007, that explains why it is most appropriate to attend, in one's practical reasoning, to those reasons that are most likely to make a difference to what you ought to do. Hence, it will often be appropriate to attend, in one's reasoning, specifically to facts about benefiting friends and family (among others).

⁸ This needs to be refined. It is presumably compatible with plausible ideals of friendship that, for instance, if you happen to be teaching a close friend and giving them a good grade (knowing this will benefit the friend considerably) you do *not* do so because it will benefit them, but because the work merits it.

⁹ Alienation (as we are using the term) essentially involves the inhibition of states of an agent. Hence only things that can inhibit states of an agent can alienate. On the assumption that ethical principles are abstract objects, the unknowable truth of an ethical theory that no-one accepts could not be alienating. (Although knowing that the ethical truth is unknowable might be.) It is individuals or collectives accepting (or believing, committing to, etc.) ethical principles that is inhibiting. An ethical theory is alienating only if it is accepted (believed, committed to, etc.) either by oneself or others. In the latter case, social acceptance of some ethical principle may be manifest in patterns of behavior, social norms, or institutions that are alienating for particular individuals – or perhaps for everyone. In the former case, by accepting an ethical theory, one is oneself inhibited from participation in some kind of normative ideal, perhaps of integrity. This marks one important difference between alienation in the current context and the more familiar context of political economy. In the latter case, the socio-economic systems that alienate are themselves the inhibitors; in the former case, ethical theories don't really alienate; acceptance of them (by oneself or others) does so.

It may also happen that, in many cases, one has *more reason* to help one's friends than strangers, just on purely impartial grounds (cf. Jackson, 1991). Take a case in which you are more motivated to provide the same resource to a friend than to a stranger in virtue of having better evidence about the contribution of the benefit to your friend's welfare. Your evidence suggests that the resource will have the same impact on their welfare. But you are justified in being more confident that the resource will make this impact on your friend than on the stranger (assume the only alternative to this positive welfare impact would be welfare neutral). This is a case in which your motivation to help your friend is stronger than your motivation to help the stranger. And, other things being equal, you will, in fact, help your friend. Continue to assume that the content of your motivation is not relevantly alienating. In this case, the *resultant* strength of your motivating reason is greater for your friend than for the stranger, for the same anticipated welfare impact.

So, there is some sense in which the Harmonious Act Consequentialist can not only advocate being motivated distinctly by facts about one's friends, but also, in some cases, more motivated by those facts than by facts about the impacts of one's actions on strangers. Is this latter fact about the resultant strength of your motive sufficient to properly recognize your friend in your practical reasoning? We suggest not. A friend might reasonably hope for more recognition than this: for their interests to play a more significant role in your practical reasoning (cf. Darwall, 1977). This is clearest when we step back to consider the rule one is following in this case: that one ought to give someone's interests more weight in proportion to one's justified confidence in successfully helping them. This entails, in many other cases, that one will give the stranger's interests more weight than one's friends whenever it happens that one's evidence about the contribution to their welfare is slightly stronger. Intuitively, following this rule in one's practical reasoning is not sufficient to properly recognize one's friend as such.

Furthermore, although the motive content of the Harmonious Act Consequentialist does include a direct reference to their friend, the agent is motivated to help their friend *in virtue of* their friendship (or their associated concern for their friend) only insofar as facts about their relationship are relevant to the consequentialist calculation (e.g., 'my friend's happiness will be increased more by my visit than by a stranger's visit'). In the agent's mind, the friendship has no normative or motivational force outside the role it plays in calculations about neutral value.

Those considerations apply to the case in which Joe is a good friend. But the alienation challenge does not merely afflict cases involving partiality—as the history of the concept of alienation in political contexts also suggests. An agent might be alienated from a perfect stranger, to whom they have no special commitments or obligations, in virtue of their failure to participate in some ideal of community. Perhaps the proper adherence to norms manifesting mutual concern or recognition involves some non-maximizing, or even some non-promotion-based, response (cf. Honneth, 1992; Scanlon, 2008; Cohen, 2009; and Scheffler, 2015). Common sense accounts of virtues such as compassion, benevolence, and sympathy, and of social roles such as citizen, neighbor, and teacher, will also often require a less abstract, less calculative, or less impartial response than the Harmonious Act Consequentialist's (cf. Brennan and Pettit's

discussion of virtues in their 1986; see also Blum's discussion of vocations and challenges to the personal/impersonal distinction in his 1993).

This introductory discussion suggests that, in doing something that benefits someone, none of the following is sufficient to avoid alienation: having the fact that you are benefitting them as part of your motivation; being distinctly motivated by the fact that you will benefit them; being saliently motivated by the fact that you are benefitting them; or being more motivated by the benefit to them than by the same benefit to another. The discussion also raises various questions. One concerns the *stringency* of your motivation: that you take your friend's interests to be more important than others in the right kind of way. One concerns the *robustness* of your motivation: that you take your friend's interests to be so important across the right range of circumstances. Another concerns the *source* of this robustness: this more modally robust orientation is based on a concern for one's friend as such. A further issue concerns the underlying *justification* for one's motivation, that it is permissible, all things considered, to be so motivated in matters concerning one's friend. Analogues of these points may also apply to actions helping strangers, perhaps also to actions in fulfilment of projects. In the remainder we will consider four kinds of strategies employed by consequentialistically-inclined theorists in response to these questions.

III. Some Alternative Consequentialist Strategies

3.1 Hybrid Theories

According to Hybrid Theories, consequentialism is not supreme. The deontic facts explained by value promotion do not entail facts about what one *just plain* ought to do.¹⁰ On some versions, a consequentialist principle fully explains the moral ought, but it is a further question how the moral ought interacts with other kinds of deontic oughts, in particular those arising from personal relationships and projects. On this view, personal relationships are beyond morality. Other versions don't distinguish a moral ought from the just plain ought. Instead they hold that morality itself has a broadly consequentialist structure, but that it permits exceptions from value maximization for acts in the context of personal relationships and other commitments (see Scheffler, 1982).

The implications for these two versions of a Hybrid Theory are different.¹¹ On the 'Distinct Oughts' view, the moral ought might be explained consequentialistically, perhaps by Act Consequentialism; it is just that morality isn't always overriding. The overall *ethical* theory includes other non-moral sources of authoritative normativity. This view is structurally similar to W.D. Ross's deontological view (1930), with impersonal benevolence as one *prima facie* duty among others. This is a case in which alienation is avoided by the overall ethical theory only by the authority of fundamental non-consequentialist principles. Furthermore, within its own jurisdiction, consequentialism still faces an alienation challenge. On the 'Partial Option' view, morality itself permits responsiveness to partial considerations. No further appeal is made to distinct moral and all things considered oughts. On this view, it is not clear whether morality is alienating at

¹⁰ For a discussion of the 'all things considered' or 'just plain' ought, see McLeod, 2001 and Maguire & Woods (op. cit.).

¹¹ Many thanks here due to Tristram McPherson.

all. However, again, it is significant that the title of the seminal defense of this position is *The Rejection of Consequentialism*.

It is also worth noting that both versions may also run into difficulties accounting for non-alienation in *non*-partiality cases. On either version of the Hybrid Theory, the agent's motive in doing something that helps a stranger will be the same as that of the regular Harmonious Act Consequentialist (assuming there are no non-moral oughts, or partial permissions, in play). As we have seen, it is an open question whether this motivational profile can fully account for the role of virtues, social roles, and civic relationships in our ethical life.

3.2 Relative Value Theories

According to Relative Value Theories, what one ought to do is explained, not by what would promote the most neutral value, but by what would promote the most 'relative value', that is, the most value relative to the agent.

There are various ways to understand the notion of 'value relative to the agent'.¹² One is to start with some idealization of the agent's desires, or with some fundamental notion of what an agent has reason to desire, or with what it would be fitting for the agent to desire (see e.g. Chappell, 2019 and Portmore, 2014). Such facts may then be used to explain what an agent ought to do and how they ought to be motivated (if motives and desires are distinct). Our central concern with such approaches is that they abandon the key consequentialist idea altogether, which is that deontic facts are explained by the agent's relation to neutral value.

An alternative starts with facts about neutral value and adds the idea of modification: these values are ranked by facts about their 'moral distance' from the agent. The function from the neutral values and their moral distance from the agent yields a set of "relative value" facts. These are then used to explain the deontic facts (for versions of this gambit, see Hurka, 2001; Maguire, 2013; Maguire, 2017; and Bader, 2016).

How would this apply to ethics of motives? Harmonious Relative Value Theories will maintain that one ought to be motivated to promote states of affairs that are relatively valuable, in proportion to their relative value. To (over)simplify matters, assume, on the one hand, that friendship calls for responsiveness just to welfare facts, and, on the other hand, that all and only welfare facts are neutrally valuable. Then the *facts* that motivate one will be the same whether one is responding just as a friend, or just as a consequentialist, or as a Harmonious Relative Value Theorist. What about the *weight* of these motives? Does the relative value of some benefit to a friend correspond to a motive of the same strength as might be demanded by some plausible normative ideal of friendship? This will turn on how particular theories characterize 'moral distance.' But this concept is designed, more or less, to yield the correct result here. So far, then, there is no obvious alienation in cases of partiality.

It is worth looking closely at the difference between the resultant motivational profile of the Harmonious Jackson-style Decision Theoretic Consequentialist and

¹² For concerns about the very idea, see Schroeder 2007.

of the Harmonious Relative Value Theorist. The latter is motivationally sensitive to more than agent-neutral facts about the promotion of welfare, relative to available evidence. In addition, they are motivationally sensitive to whatever it is that accounts for 'moral distance.' Suppose, for a moment, that one is 'morally closer' to one patient than another to the extent that one actually cares more about the one than the other. Then the strength of one's motivation to help someone one cares more about than another will be explained partly by the facts about the neutral value of welfare and also partly by the fact that one cares more about the one than the other. That seems like just the kind of motivational sensitivity that might avoid alienation.

Furthermore, in this case, as contrasted, in particular, with the 'Distinct Oughts' version of a Hybrid Theory, alienation is not avoided by one principle *in spite of* the alienating consequentialist moral principle. Rather the consequentialist principle is playing a partial role in the full explanation of the all things considered ought, alongside moral distance. And, indeed, it seems proper for a friend to be responsive to the (neutrally valuable) facts about the impacts on their friend's welfare of this or that intervention.

But there is a similar concern here about what in this ethical theory secures non-alienation. The concern is that it is not exactly the *consequentialist* principle that gets the motive weight right, but whatever explains one's 'moral distance' from the patient. The notion of moral distance is not coming from the consequentialism but from some other principle: perhaps strength of care, as in our example, or perhaps simple egoism. Whatever plays this role has fundamental normative authority.

One possibility for the Relative Value Theorist is to offer a consequentialist account of 'moral distance' (cf. Maguire, 2016). For it is rather implausible that actual facts about an agent's psychology are fully adequate to play the role of a normatively authoritative modifier. One often fails to care as one should. Perhaps, then, moral distance can be explained by what one *ought* to care about; this 'ought' might, in turn, admit of a simple consequentialist account. It is hard to believe that, case by case, more value will be promoted by helping the near and dear. But it is more plausible – particularly in more egalitarian worlds – that having a caring profile according to which one cares more about one's near and dear than strangers will promote more value than a caring profile that ranks everyone equally (cf. Maguire, 2016).

We are doubtful whether such a position could yield fully satisfying extensional results concerning actions and motives. But even if so, this position strikes us as unnecessarily baroque. It is the conjunction of an Act Consequentialist theory of the existence of reasons with a combination of Act Consequentialist and Levelled Consequentialist theory of the weight of reasons. At this point, a consequentialist would be advised to opt instead for a full-blooded levelled consequentialism, against which we will shortly present a range of challenges.

3.3 Global Consequentialism

Global Consequentialism is the universalized version of direct consequentialism.¹³ According to Global Consequentialism, the most value-promoting alternative for every given 'focal point' (actions, motives, forms of government, roof colors, etc.) ought to obtain.

Global Consequentialism is often presented as a solution to the problems that face Harmonious Act Consequentialism. The latter gives the wrong answer to the question of what ought to motivate an agent in interpersonal contexts (both when the agent is interacting with the near and dear and with strangers) and, plausibly, in certain intrapersonal contexts, as when an agent is pursuing a ground project (on which more shortly). But in many cases, having the 'non-alienated' motive in such contexts promotes more value than not. This permits a conjecture: direct consequentialism about motives might have a more intuitively attractive set of deontic prescriptions for motives (cf. Sidgwick, 1874 and Adams, 1976). The Global Consequentialist is unwilling to relinquish their direct assessment of acts, and instead relinquishes the idea that there is any direct normative relationship between the assessment of acts and the assessment of motives. For that matter, it also relinquishes any direct normative relationship between ways of reasoning and actions, between commitments and ways of reasoning, etc.

As far as alienation is concerned, there are various problems.¹⁴ The first is this. It may be true that the motives, commitments, modes of reasoning, etc. advocated by some ideal of friendship are consequentialistically better than those advocated by Harmonious Act Consequentialism. But it does not follow that the former are the *best* available – not even taking into account the disvalue, if there is any, of alienation itself. Presumably it will often be consequentialistically better to be motivated to abandon our friends and family in a cocaine-fueled pursuit of hedge-fund-managed millions, to be gifted to the poor.

The second problem is this. Take a case in which, according to Global Consequentialism, you ought to be most motivated to x and yet in which you ought not x . Lots of familiar kinds of partiality cases will be like this: you ought to be motivated to help your friend or your child, but you ought to help the slightly needier stranger. The problem is that, in many such cases, the very normative ideal that provides support for the claim that you ought to be motivated to help your friend rather than the stranger also provides support for the claim that you ought to help your friend rather than the stranger; and that you ought to help your friend *out of* your greater concern for your friend. Normative ideals can call for partial rather than impartial behavior. But since Global Consequentialism entails Act Consequentialism, it entails that one ought always to do whatever is impartially best.

A closely related problem concerns the intrapersonal ideal of integrity. According to Global Consequentialism, the fact that one ought to have some commitment that would make it fitting to be motivated to x in S is explanatorily irrelevant to whether one ought to be so motivated, since being committed and being motivated

¹³ For general discussions and overviews of Global Consequentialism, see Driver, 2014 and Ord, ms. For defenses of the view, see Parfit, 1984; Kagan, 2000; Pettit and Smith, 2000; and Feldman, 1993.

¹⁴ Special thanks to Austen McDougal for discussion.

are distinct focal points.¹⁵ Suppose that one ought not to be so motivated. A suitably apprised Global Consequentialist will find themselves struggling to satisfy these two directives. If they have the motives they ought to have, they will fail to have the motives that fit their commitments – in contravention of a plausible normative ideal. Commitments can survive occasional exceptions. But it is not so clear that they can survive the sober, clear-headed judgement that they have no motivational significance, now or ever. So, having the Global Consequentialist motive also undermines one's claim to having the relevant commitment at all.

One way to sum this up is to point out that non-alienation relative to a range of prominent normative ideals calls for *multi-focal fittingness*: for an agent to be oriented towards another in their actions, motives, affect, and ways of thinking, where these different states are all normatively integrated at a time and across time. In disaggregating these different responses, Global Consequentialism gives precisely the wrong kind of prescription.

IV. Levelled Consequentialism

4.1 Characterization of Levelled Consequentialism

The discussions so far have variously pointed in the direction of a version of consequentialism that directly assesses higher-level psychological states, which in turn yield more specific agential attitudes. Call such theories Levelled Consequentialisms (see e.g. Rawls, 1955; Railton, 1984; Pettit and Brennan, 1986; Tännsjö, 1995; Mason, 1998; Mason, 1999; Tedesco, 2006; and Norcross, 1997). Examples of such higher-level states include commitments, values, identities, decision procedures, roles, caring profiles, character traits, and, simply, dispositions. These psychological structures are robust in the sense that they tend to persist over time and possibilities. They are higher-level in the sense that they explain – causally, by making fitting, or both – a range of more specific, lower-level states such as motives, thought patterns, and actions. The Levelled Consequentialist has a 'two-tier' psychological structure, and these two tiers are potentially responsive to different kinds of justifications. The higher-level states generate, coordinate, and support a set of motives, thought patterns, and/or behavioral tendencies. The higher-level states are responsive, in their turn, to consequentialist considerations.

There are both indirect and global versions of Levelled Consequentialism, differing as to whether they maintain that one ought to act in response to one's higher-order states, or whether one ought nonetheless to act in whatever way would be best.¹⁶ Global Consequentialism disaggregates the two, maintaining that you ought to do whatever is best and that you ought to be committed however is best, even though the two might conflict (Hare, 1981; Railton, 1984; Pettit & Brennan, 1986). Indirect Levelled Consequentialism maintains that you ought to do whatever you ought to be committed to (or disposed to, or what have you) (Rawls, 1955; Tedesco, 2006; Mason, 1998).

¹⁵ An interesting and concerning feature of Global Consequentialism, which has not been much remarked upon (though see Louise 2006, 82-84), is that the putatively non-ethical individuation of focal points comes to carry a great deal of ethical weight.

¹⁶ Pity the fool who attempts to read this distinction back into the literature.

Notice that indirect Levelled Consequentialisms need not be ‘self-effacing’.¹⁷ To see this, stick with Levelled Consequentialisms that involve principles of reasoning. Such theories maintain that one ought to employ some non-consequentialist decision procedure in certain first-order contexts. But they also maintain that the ethical facts about actions are themselves responsive to such non-consequentialist considerations in those contexts. Such theories do tell one to employ a consequentialist decision procedure in higher-level matters, but again the ethical facts are responsive to consequentialism on that level. So, there is congruence at both levels between the criterion of rightness and the right decision-making procedure. But, as before, this comes at the ‘cost’ of relinquishing explanatory autonomy to non-consequentialist considerations.

Levelled Consequentialisms – especially indirect Levelled Consequentialisms – strive to take seriously the idea that some of the most important goods in life consist in proper participation in certain normative ideals, such as those of familial love or and caring friendships. These theories aim to preserve the fact that such participation can require not just behavior but motives, affect, particular ways of reasoning, and, in general, commitments to the activity for its own sake. The two-tier structure attempts to uphold the autonomy of standards of participation, while restricting fundamental authority to the consequentialist assessment of participation itself.

We wish to raise three kinds of concerns about the levelled approach. The first question concerns its feasibility: can the distinction between levels be upheld? A second question concerns its applicability: does the alienation objection also apply at the higher level? What conditions are acceptable on the commitment of a friend, citizen, etc.? The third question concerns integrity: is the Levelled Consequentialist thereby alienated from their own particular commitments, or types of commitments, or their relation to themselves as an agent with commitments?

4.2.1 The Collapse Objection

According to the Collapse Objection, the distinction between two levels does not enable the agent to avoid alienation, for the fact that the higher level admits of consequentialist justification unavoidably infects the standing or nature of the lower-level states. To take a specific version of the objection, one is not *really* doing something for one’s friend because of one’s friendship, but because the friendship is supported by the most value promoting disposition, commitment, or whatever.

We are optimistic that this concern can be avoided. The structuring of practical thought into distinct levels, spheres, domains, etc., seems to us a deep and unavoidable fact of life. As Rawls, 1955 argued, we frequently sustain our commitment to one set of rules, which provide their own set of justifications, on the basis of our commitment to some other kind of justification. For instance, we might accept a practice of punishment for wrongdoing because of its utilitarian advantages, but punish independently of the utilitarian advantages of doing so. To take another example, we are fluent in transitions from “I” reasoning to “we” reasoning in everyday life; we can also reason differently *qua* teacher, *qua* parent, and *qua* friend. Or to take a different example from Dean Cocking and Justin Oakley, “someone who is prepared to continue serving a patient’s health only

¹⁷ See fn. 7 above for the distinction between indirect and self-effacing.

provided he earns enough money to subsidize his passion for the opera could still be a very admirable doctor" (2001, 87). There is a mixture of empirical, theoretical, and normative issues here that we cannot sort out just now (for recent work on associated empirical issues see Seligman et al., 2017). Suffice it to say that interesting work is required to be done in distinguishing the cases in which levels do collapse, of which Rawls' 'summary rules' are an example, from those in which they do not, of which Rawls's 'practice rules' are an example.

4.2.2 The Higher-Level Alienation Objection

Earlier we encountered the plausible conjecture that non-alienation requires not just particular motives, but also that those motives have a particular source: that one is motivated to do things out of the right kind of commitment or other higher-level state. Here is a related concern. It might be argued that certain kinds of reasons for having friendly commitments can themselves undermine the status of the relationship as a friendship (Kapur 1991 and Cocking and Oakley, 1995). The issue here doesn't concern the status of the first-order motives. It may be that the collapse can be avoided. But it is charged that the higher-order motives themselves are sufficient to vitiate the normative ideal.

To take an important example, Cocking and Oakley argue that the ideal of friendship dictates certain conditions under which one would enter into and leave a friendship – the 'governing conditions' of one's commitment to the relationship. They argue that the responsiveness of higher-level states to consequentialist considerations would mean that one would fail to meet these 'governing conditions' for one important ideal of friendship, and hence that one would fail to be a friend at all. Indeed, they maintain that the governing conditions are 'most significantly responsible for the problem of alienation' (1995, 111).

Presumably, certain kinds of conditionality are compatible with friendship. You might be disposed to relinquish your commitment if you discover that your friend is a serial killer, or if they have an affair with your spouse, or simply if your relationship ceases to be enjoyable for you both. That conditionality doesn't interfere with the quality of your relationship if there have been no murders or affairs. But Cocking and Oakley argue that straightforward conditionality of a consequentialist kind would be incompatible with friendship.

By way of analogy, they discuss the example of an ambitious philosopher befriending a famous professor for professional advancement. However, the ambitious philosopher has successfully ensured that only their commitment to the friendship is responsive to their professional goals; their motives are not counterfactually responsive in this way. The individual actions of the ambitious philosopher are not motivated by professional ambition, but the philosopher is disposed to relinquish the relationship if it no longer serves their ambition.

We agree that there is alienation in the professor case, relative to a familiar normative ideal of friendship. However, the alienation feels a little different to us. Assume the Collapse Objection can be avoided, and that the ambitious philosopher really does *care* about the famous professor. Cocking and Oakley allow that this commitment may not be dynamically responsive to the philosopher's judgement that the friendship no longer serves their ambitious interests. The philosopher might need to work hard to get themselves to stop

caring so much about the famous professor, perhaps to focus on networking with the deans instead. Granted, something is wrong with their relationship with the professor. But something is right about it too. The professor would have been *more* alienated had the philosopher's ambition led them to try to gainsay their opinions at every turn, for instance.

An option for the consequentialist at this point is to add another level. Perhaps it is not individual relationships themselves that are assessed for value maximization, but the governing structures that coordinate and support individual relationships (Mason, 1998; Tedesco, 2006). But in reply, one might modify Cocking and Oakley's objection and argue that it is alienating for a Levelled Consequentialist's friendships to be conditional upon a higher-level governing structure, which one accepts if and only if it is appropriately value promoting. But at this point, while there is clearly something distasteful about the ambitious philosophers' character, it is less clear that the ambitious philosopher would be alienated from the famous professor. Imagine if all these facts were common knowledge, and the famous professor accepted the higher-level governing structure on the same basis. Perhaps the famous professor is friends with the ambitious philosopher only on the basis of a general commitment to having others think they are generous and gregarious. But they both really do care about each other. This pair sound vicious but not necessarily alienated – or as alienated – from each other.

When we turn from conditionality on egoistic considerations, to conditionality on consequentialist considerations, things are even murkier. For it is plausible enough that *some* higher-order conditionality on purely extrinsic moral matters is permitted by familiar ideals of friendship. Suppose that you are an epidemiologist with specialized training in the treatment of some rare disease in the days before long distance communication. You might be friends with someone even though you both know that you might have to move overseas indefinitely and at short notice. There would be certain kinds of investments you would both be less likely to make in the friendship. You won't plan to buy a house together, for instance. So, there may be degrees of commitment that you are prevented from participating in, and you might be alienated from ideals involving the prospect of such a degree of commitment. But you might have a valuable committed friendship all the same, with most of your first-order responses intact. You see each other a couple of times a month to enjoy shared activities; you help each other out occasionally; you check in when one or the other has something to celebrate or commiserate about; etc.¹⁸ As this conditionality ascends to higher levels in one's commitment structure, it likewise sounds less problematic. Perhaps one is disposed to stop having any close friends at all, if circumstances change enough that having friends can no longer be justified on consequentialist grounds.

However, notice that this correlates with consequentialism itself playing an increasingly diminished role in the overall organization of one's life. As the consequentialist assessment ascends to higher levels, one is responsive to more and more commitments, each with its own non-consequentialist authority.

4.2.3 The Moral Self Objection

¹⁸ It is easy to construct examples in which the friendship was initiated for the wrong kind of reasons, as well.

A further question remains, namely whether one will be alienated from *oneself* by one's commitment to Levelled Consequentialism.¹⁹ A certain amount of alienation from our actual commitments is a necessary condition on being able to reflect critically on one's own, and our collective, life (cf. Railton, 1986). It is not problematic for an ethical theory to require the egoist or militant racist or perhaps even the masochist to modify or abandon a commitment, project, role, relationship, or personality trait. And perhaps the true ethical theory will require most of us to adopt more of the sorts of commitments consequentialists talk about, for instance to donating more to organizations helping people in extreme poverty, resisting the factory farming system, and caring more about future generations, all at a trade-off to our current lifestyles and resource allocations. But all this leaves open the questions of whether the true ethical theory subjects our identities, ground projects, etc. to *consequentialist* assessment, and whether an ethical theory that does so is alienating. Insofar as consequentialism evaluates one's higher-order states just in terms of the neutral value of their total consequences, people will likely be unable to honor a range of deeply held, and intuitively reasonable, ethical and non-ethical commitments. The key question here is: What standards of ethical assessment can one accept without somehow alienating oneself?

The challenge for consequentialism is perhaps clearest in cases where value differences are marginal. Consider a quantitatively-skilled public policy professional working on US education who, after some careful research, expected value calculations, and considerations of replaceability in the job market concludes that the expected impact of their career would be slightly higher if they switched to hedge fund management and donated much of their earnings to the best non-profit working on education. Imagine further that this person is fascinated by public policy and political science, but bored by finance and the stock market. It seems alienating for an ethical theory to require this person to give up their fulfilling career in policy for a boring, career in finance simply because of a marginally higher expected impact. To take another example, imagine someone whose chief hobby in life is chess. They realize that if they gave up chess and devoted the time and energy to gardening instead, they would generate more neutral value, since people could aesthetically appreciate the resulting garden. Again, it seems alienating for an ethical theory to require this person to renounce chess in favor of gardening, simply because the change of project would promote somewhat more neutral value.

As pointed out in the discussion of Global Consequentialism, it doesn't follow from the fact that a non-alienating set of commitments would add value that the non-alienating set would be consequentialistically *best*; nor that even if it were, consequentialism provides the intuitively correct justification for the commitments. Likewise, it does not follow from the evident fact that many peoples' commitments are morally dubious, and even that they are insufficiently sensitive to facts about neutral value promotion, that the proper standard of critical assessment is some version of Supreme Consequentialism. Particular commitments may involve sensitivity to neutral value in their own way. For instance, caring about a friend involves some sensitivity to what is good for them, which is (at least conditionally) of value. Likewise committing to being a good doctor involves positive sensitivity to values. And one might have a larger

¹⁹ Cf. Williams, 1973; Williams, 1981a; and Williams, 1981b.

commitment to regulating one's commitments in part by their social usefulness, measured roughly consequentialistically. But this falls far short of finding it acceptable to assess one's commitments wholesale just in terms of their consequential value.

One final point. For all we have said, someone, or some society, might come to accept some version of Supreme Consequentialism on the basis of authentic ethical reflection. They might also accept compatible normative ideals concerning the nature of the self, and of society, and of the very enterprise of morality. But this kind of coherence is not sufficient for non-alienation. On the contrary, this would just make them *more* alienated relative to alternative and more attractive ideals (see Cohen 1968, 215, on "ersatz humanity").

V. Conclusion

We considered four consequentialist strategies for avoiding alienation objections of various kinds. None is fully satisfying.

One option – Hybrid Theory – is to restrict one's consequentialism to a subordinated role in one's overall theory of what one ought to do. One worry with this view is whether the residual consequentialism will itself be alienating, offering too abstract a standard even in impersonal cases. The more general concern is that the overall theory of what one ought to do is not consequentialist, and alienation is avoided just by the non-consequentialist part. This seems less a victory for consequentialism than a defeat.

Another option – Relative Value Theory – is to mix one's consequentialist assessment of options with an agent-relative ranking, using the resulting agent-relative ranking to explain facts about what one ought to do. The central worry with this approach is that, again, insofar as alienation is avoided, it is by the non-consequentialist constituent in the proposal, namely the grounds for whatever yields the agent-relative ranking.

A third option – Global Consequentialism – is to disaggregate one's assessment of motives and actions. Plausibly enough, an impartial set of motives will promote less value than a partial set of motives; partial motives can then be advocated on the grounds of neutral value promotion. Problems will persist when non-alienation intuitively requires not just partial motives but partial actions. Further problems arise upon consideration of the agent in conformity with Global Consequentialism – their values, commitments, motives, and actions all askew.

In these cases, however, progress seems to be made in ethical theory, if not exactly towards a consequentialist ethical theory, for we have a better understanding of how we might integrate values and commitments with different kinds of ethical assessment.

The fourth option – Levelled Consequentialism – advocates compliance with higher-level states, such as one's commitments. It is plausible that a consequentialist assessment of such higher-level states will advocate commitments to oneself and others in ways that involve less alienation, at least, than traditional Act Consequentialism. Or perhaps one can restrict

consequentialism to even higher states that govern one's selection of commitments, or perhaps of social, political, and economic institutions and practices within which one lives one's life. At each greater height, consequentialism cedes more normative authority to non-consequentialist rules and institutions. But it seems also to be alienating to make certain trade-offs, and to make or reject certain commitments, on purely consequentialist grounds, even at these lofty theoretical heights.

It remains an open question whether consequentialism will play some more limited explanatory role in a fully satisfying ethical theory, or whether the considerations that speak in its favor – facts about proper sensitivity to value, the appropriateness of dominance reasoning or maximization in different domains, and so on – will be fundamentally explained in other ways, which may be more straightforwardly reconciled with our normative ideals.

References

- Adams, R. M. (1976). "Motive Utilitarianism." *The Journal of Philosophy* 73, no. 14: 467–81. <https://doi.org/10.2307/2025783>.
- Arpaly, Nomy (2011). *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.
- Bader, Ralf (2016). "Conditions, Modifiers, and Holism." In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 27-55. New York: Oxford University Press.
- Blum, Lawrence (1993). "Vocation, Friendship, and Community: Limitations of the Personal-Impersonal Framework." In *Identity, Character, and Morality*, edited by Owen Flanagan, 173-98. Cambridge: MIT Press.
- Brudney, Daniel (2014). "The Young Marx and the Middle-Aged Rawls." In *A Companion to Rawls*, edited by Jon Mandle and David Reidy, 450–71. Malden: Wiley Blackwell.
- Cocking, Dean, and Justin Oakley (1995). "Indirect Consequentialism, Friendship, and the Problem of Alienation." *Ethics* 106, no. 1: 86–111. <https://doi.org/10.1086/293779>.
- Cohen, G.A. (1968). "Bourgeois and Proletarians." *Journal of the History of Ideas* 29, no. 2: 211-230. <https://doi.org/10.2307/2708577>.
- Cohen, G.A. (2009). *Why Not Socialism?* Princeton: Princeton University Press.
- Cox, Damian (2005). "Integrity, Commitment, and Indirect Consequentialism." *The Journal of Value Inquiry* 39, no. 1: 61–73. <https://doi.org/10.1007/s10790-006-1571-7>.
- Crisp, Roger (1992). "Utilitarianism and the Life of Virtue." *The Philosophical Quarterly* 42, no. 167: 139-60. <https://doi.org/10.2307/2220212>.
- Dancy, Jonathan. (2019). *Practical Shape*. Oxford: Oxford University Press.
- Darwall, Stephen (1977). "Two Kinds of Respect." *Ethics* 88, no. 1: 36-49. <https://doi.org/10.1086/292054>.
- Driver, Julia (2014). "Global Utilitarianism." In *The Cambridge Companion to Utilitarianism*, edited by Ben Eggleston and Dale E. Miller, 166–76. New York: Cambridge University Press.
- Feldman, Fred (1993). "On the Consistency of Act- and Motive-Utilitarianism: A Reply to Robert Adams." *Philosophical Studies* 70, no. 2: 201–12. <https://doi.org/10.1007/bf00989590>.
- Hare, R.M. (1981). *Moral Thinking: Its Levels, Method, and Point*. Oxford: Oxford University Press.
- Honneth, Axel. (1992). *The Struggle for Recognition*. Cambridge: MIT Press.
- Hooker, Brad (2000). *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Clarendon Press.
- Hurka, Thomas (2001). *Virtue, Vice, and Value*. New York: Oxford University Press.
- Jackson, Frank (1991). "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101, no. 3: 461–82. <https://doi.org/10.1086/293312>.
- Jaeggi, Rahel (2014). *Alienation*. Translated by Frederick Neuhouser and Alan E. Smith. Edited by Frederick Neuhouser. New York: Columbia University Press.
- Kagan, Shelly (2000). "Evaluative Focal Points." In *Morality, Rules and Consequences: A Critical Reader*, edited by Brad Hooker, Elinor Mason, and Dale E. Miller, 134–55. Edinburgh: Rowman & Littlefield.

- Kain, Philip (1982). *Schiller, Hegel, and Marx*. Kingston and Montreal: McGill Queens University Press.
- Kapur, Neera Badhwar (1991). "Why It Is Wrong to Be Always Guided by the Best: Consequentialism and Friendship." *Ethics* 101, no. 3: 483–504. <https://doi.org/10.1086/293313>.
- Kaufman, Walter (1970). Introduction to *Alienation*, by Richard Schacht. Garden City, New York: Doubleday.
- Leopold, David (2018). "Alienation." In *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/fall2018/entries/alienation>.
- Louise, Jennie (2006). "Right Motive, Wrong Action: Direct Consequentialism and Evaluative Conflict." *Ethical Theory and Moral Practice* 9, no. 1: 65–85. <https://doi.org/10.1007/s10677-005-9000-8>.
- MacIntyre, Alasdair C. (1984). *After Virtue: A Study in Moral Theory*. Notre Dame: University of Notre Dame Press.
- Maguire, Barry (2015). "Grounding the Autonomy of Ethics." In *Oxford Studies in Metaethics: Volume 10*, edited by Russ Shafer-Landau, 188–215. Oxford: Oxford University Press.
- Maguire, Barry (2017). "Love in the Time of Consequentialism." *Nous* 51, no. 4: 686–712. <https://doi.org/10.1111/nous.12169>.
- Maguire, Barry (2016). "The Value-Based Theory of Reasons." *Ergo* 9, no. 3: 233–62. <http://dx.doi.org/10.3998/ergo.12405314.0003.009>.
- Maguire, Barry (2013). "Values, Reasons, and Ought." PhD dissertation, Princeton University.
- Maguire, Barry and Jack Woods (fc.). "The Game of Belief." *The Philosophical Review*.
- Markovits, Julia (2010). "Acting for the Right Reasons." *The Philosophical Review* 119, no. 2: 201–42. <https://doi.org/10.1215/00318108-2009-037>.
- Mason, Elinor (2002). "Against Blameless Wrongdoing." *Ethical Theory and Moral Practice* 5, no. 3: 287–303. <https://doi.org/10.1023/A:1019671210369>.
- Mason, Elinor (1998). "Can an Indirect Consequentialist Be a Real Friend?" *Ethics* 108, no. 2: 386–93. <https://doi.org/10.1086/233810>.
- Mason, Elinor (1999). "Do Consequentialists Have One Thought Too Many?" *Ethical Theory and Moral Practice* 2, no. 3: 243–61. <https://doi.org/10.1023/A:1009998927955>.
- McLeod, Owen (2001). "Just Plain "Ought"." *The Journal of Ethics* 5, no. 4: 269–91. <https://doi.org/10.1023/A:1013934513554>.
- Norcross, Alastair (1997). "Consequentialism and Commitment." *Pacific Philosophical Quarterly* 78, no. 4: 380–403. <https://doi.org/10.1111/1468-0114.00045>.
- Oakley, Justin, and Dean Cocking (2001). *Virtue Ethics and Professional Roles*. Cambridge: Cambridge University Press.
- Ord, Toby (unpublished manuscript). "How to Be a Consequentialist about Everything." Available at <http://www.amirrorclear.net/academic/papers/everything.pdf>.
- Parfit, Derek (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Pettit, Philip and Geoffrey Brennan (1986). "Restrictive Consequentialism." *Australasian Journal of Philosophy* 64, no. 4: 438–55. <https://doi.org/10.1080/00048408612342631>.
- Pettit, Philip and Michael Smith (2000). "Global Consequentialism." In *Morality, Rules and Consequences: A Critical Reader*, edited by Brad Hooker, Elinor Mason, and Dale E. Miller, 121–33. Edinburgh: Rowman & Littlefield.

- Portmore, Douglas W. (2014). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford: Oxford University Press.
- Railton, Peter (1984). "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13, no. 2: 134–71. <https://www.jstor.org/stable/2265273>.
- Rawls, John. (1955). "Two Concepts of Rules." *The Philosophical Review* 64, no. 1: 3–32. <https://doi.org/10.2307/2182230>.
- Rawls, John (1955). "Two Concepts of Rules." *The Philosophical Review* 64, no. 1: 3–32. <https://doi.org/10.2307/2182230>.
- Raz, Joseph (1986). *The Morality of Freedom*. Oxford: Oxford University Press.
- Ross, W.D. (1930). *The Right and the Good*. Oxford: Clarendon Press.
- Scanlon, T.M. (2008). *Moral Dimensions*. Cambridge: Harvard University Press.
- Schacht, Richard (1970). *Alienation*. Garden City, New York: Doubleday.
- Scheffler, Samuel (2010). *Equality and Tradition*. Oxford: Oxford University Press.
- Scheffler, Samuel (2015). "The Practice of Equality." In *Social Equality: Essays on What it Means to be Equals*, edited by Carina Fourie, Fabian Schuppert, and Ivo Wallimann-Helmer, 21–44. Oxford: Oxford University Press.
- Scheffler, Samuel (1982). *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Clarendon Press.
- Schroeder, Mark (2007). "Teleology, Agent-Relative Value, and 'Good.'" *Ethics* 117, no. 2: 265–95. <https://doi.org/10.1086/511662>.
- Seligman, Martin et al. (2017). *Homo Prospectus*. Oxford: Oxford University Press.
- Sidgwick, Henry (1981). *The Methods of Ethics*. Indianapolis: Hackett Publishing Co. First published 1874.
- Stocker, Michael (1976). "The Schizophrenia of Modern Ethical Theories." *Journal of Philosophy* 73, no. 14: 453–66. <https://doi.org/10.2307/2025782>.
- Tännsjö, Torbjörn (1995). "Blameless Wrongdoing." *Ethics* 106, no. 1: 120–27. <https://doi.org/10.1086/293781>.
- Tedesco, Matthew (2006). "Indirect Consequentialism, Suboptimality, And Friendship." *Pacific Philosophical Quarterly* 87, no. 4: 567–77. <https://doi.org/10.1111/j.1468-0114.2006.00275.x>.
- Way, Jonathan (2017). "Reasons as Premises of Good Reasoning." *Pacific Philosophical Quarterly* 98, no. 2: 251–270. <https://doi.org/10.1111/papq.12135>.
- Whiting, Daniel (2017). "Against Second-Order Reasons." *Nous* 51, no. 2: 398–420. <https://doi.org/10.1111/nous.12138>.
- Wilcox, William H. (1987). "Egoists, Consequentialists, and Their Friends." *Philosophy & Public Affairs* 16, no. 1: 73–84. <http://www.jstor.org/stable/2265206>.
- Williams, Bernard (1973). "A Critique of Utilitarianism." In *Utilitarianism: For and Against*, by Bernard Williams and J.J.C. Smart, 77–150. New York: Cambridge University Press.
- Williams, Bernard (1985). *Ethics and the Limits of Philosophy*. Cambridge: Harvard University Press.
- Williams, Bernard (1981a). "Persons, Character, and Morality." In *Moral Luck: Philosophical Papers, 1973–1980*, by Bernard Williams, 1–19. Cambridge: Cambridge University Press.
- Williams, Bernard (1981b). "Utilitarianism and Moral Self-Indulgence." In *Moral Luck: Philosophical Papers, 1973–1980*, by Bernard Williams, 40–53. Cambridge: Cambridge University Press.
- Wood, Allen (1999). *Karl Marx*. London: Routledge.