

How to be a consequentialist about everything

Toby Ord*

Over the last few decades, there has been an increasing interest in *global consequentialism*. Where act-consequentialism assesses acts in terms of their consequences, global consequentialism goes much further, assessing acts, rules, motives — and everything else — in terms of the relevant consequences. Compared to act-consequentialism it offers a number of advantages: it is more expressive, it is a simpler theory, and it captures some of the benefits of rule-consequentialism without the corresponding drawbacks. In this paper, I explore the four different approaches to global consequentialism made by Parfit, Pettit and Smith, Kagan, and Feldman. I break these up into their constituent components, demonstrating the space of possible global consequentialist theories, and I present two new theories within this space.

1. Evaluative Focal Points

Acts and rules are two examples of what Shelly Kagan calls *evaluative focal points*.¹ They are the evaluands that receive the direct consequentialist evaluation in act-consequentialism and rule-consequentialism respectively. For the purposes of this essay, we shall take act-consequentialism to be specified by the following principle:

Act-consequentialist criterion of rightness for acts

an act is right iff it would lead to a better outcome than would any available alternative act.

Rule consequentialism is more complicated, due to its indirect evaluation of acts and a number of heated questions concerning the role in which we are to consider the rules. For the purposes of this article, we shall define it by the following principles:

Rule-consequentialist criterion of the best set of rules

r is the best set of rules iff the widespread adoption of r would lead to a better outcome than would the widespread adoption of any available alternative set of rules.

Rule-consequentialist criterion of rightness for acts

an act is right iff it complies with the best set of rules.

However, acts and rules are not the only possible targets for evaluation in consequentialist theories. For example, Robert Adams has advocated a form of motive-consequentialism in which an agent's motives are directly assessed in terms

* [Acknowledgements].

¹ Kagan (2000).

of the good they would lead to, while the agent's acts are either left unassessed, or considered to be right when they spring from the best motives.²

Adams is not alone in wanting to directly apply a consequentialist criterion of rightness to a person's motives³. Indeed Bentham, Mill and Sidgwick all made claims of this type. For example, Bentham's *Introduction to the principles and morals of legislation* states:

'With respect to goodness and badness, as it is with every thing else that is not itself either pain or pleasure, so is it with motives. If they are good or bad, it is only on account of their effects: good, on account of their tendency to produce pleasure, or avert pain: bad, on account of their tendency to produce pain, or avert pleasure.'⁴

In his *System of Logic*, Mill made a similar claim, but in terms of a person's character:

'I fully admit that this is true: that the cultivation of an ideal nobleness of will and conduct, should be to individual human beings an end, to which the specific pursuit either of their own happiness or of that of others (except so far as included in that idea) should, in any case of conflict, give way. But I hold that the very question, what constitutes this elevation of character, is itself to be decided by a reference to happiness as the standard. The character itself should be, to the individual, a paramount end, simply because the existence of this ideal nobleness of character, or of a near approach to it, in any abundance, would go further than all things else towards making human life happy.'⁵

In his *Methods of ethics* Sidgwick wrote:

'the doctrine that Universal Happiness is the ultimate *standard* must not be understood to imply that Universal Benevolence is the only right or always best *motive* of action. For, as we have before observed, it is not necessary that the end which gives the criterion of rightness should always be the end at which we consciously aim: and if experience shows that the general happiness will be more satisfactorily attained if men frequently act from other motives than pure universal philanthropy, it is obvious that these other motives are reasonably to be preferred on Utilitarian principles.'⁶

² Adams (1976). Adams is unclear as to how motive-consequentialism assesses acts, although he does argue that motive-consequentialism is inconsistent with act-consequentialism.

³ Authors differ in their use of the term 'motives', sometimes specifying a meaning and often leaving it undefined. For the present purposes we can adopt Derek Parfit's definition of a person's motives as their desires and dispositions.

⁴ Bentham (1789), p. 102.

⁵ Mill (1843), Bk VI, Ch 12, §7.

⁶ Sidgwick (1907), p. 413. The emphasis is his.

Like Bentham, Sidgwick also suggested that this assessment of motives in terms of their consequences could be extended to other objects of evaluation:

‘While yet if we ask for a final criterion of the comparative value of the different objects of men’s enthusiastic pursuit, and of the limits within which each may legitimately engross the attention of mankind, we shall none the less conceive it to depend upon the degree in which they respectively conduce to Happiness.’⁷

In addition to acts, rules, and motives, there are many other natural evaluands for moral theories. For example, we may want to assess norms, character traits, decision procedures, institutions, and even entire lives. For each of these evaluands, one could construct a consequentialist theory in the mould of act-, rule-, or motive-consequentialism: directly assessing the one central evaluand and either providing some form of indirect assessment, or else remaining silent regarding each of the others.

Alternatively, we could directly assess all of these evaluands in terms of the good that they would lead to. This approach seems more in keeping with the above quotations by the classical utilitarians and has been advocated explicitly by Derek Parfit, Fred Feldman, Shelly Kagan, Philip Pettit and Michael Smith.⁸ It may also be implicitly advocated by many act-consequentialists, as it is not in conflict with their criterion of the rightness of acts. Instead, it can be seen as an extension of the minimalist act-consequentialism which I defined above, allowing the assessment of new evaluands about which the theory was formerly silent.

The advocates of this approach do not limit themselves to the list of evaluands above. Instead, they seek to take the approach to its natural conclusion and assess *everything* directly in terms of its consequences. I shall thus follow Pettit and Smith in calling this approach *global consequentialism*. Even if we set aside questions about axiology (i.e. the goodness of individual outcomes), there are many ways in which one could develop a global consequentialist theory. I will use the term to describe the general class of theories which apply a consequentialist criterion to everything, rather than to designate any specific theory within this class.

The potential advantages of global consequentialism are many. The addition of new evaluands increases the expressivity of consequentialism. It allows consequentialists to directly answer questions about the best motives, the best system of government, or the best way to decide what to do. This also allows consequentialists to bring consequentialism into the strongholds of the deontologists and virtue ethicists: it can explain the importance of character and of rules of conduct, going so far as to show how acclaimed virtues and rules have systematically led to good consequences, and even daring to suggest ways in which they might be improved.

⁷ Sidgwick (1907), p. 405.

⁸ Parfit (1984), pp. 24–28, Kagan (2000), Pettit and Smith (2000), Feldman (1986, 1993).

It promises to do all of this with few, if any, drawbacks. Unlike rule- or motive-consequentialism, it never departs from the core consequentialist ideal of assessing things by their consequences. By encompassing all evaluands, rather than just the most prominent ones, it maximizes its expressive power while remaining non-arbitrary. Indeed, some systems of global consequentialism promise to be simpler than act-consequentialism itself, for by allowing everything to be morally assessed, they no longer need an associated theory of acts.

2. Three rough accounts of global consequentialism

2.1 The three accounts

Though Parfit, Kagan, Pettit and Smith all advocate forms of global consequentialism, they do so only in general terms: defining their positions in opposition to act- and rule-consequentialism, but without filling in many details.⁹ While their main points are clear, there are nuances on which they do not elaborate. As their explanations of global consequentialism are also very concise, it will be best to simply quote them each in full rather than trying to summarise.

In *Reasons and persons*, Parfit introduces his version of consequentialism:

‘There are different versions of *Consequentialism*, or *C*. *C*’s central claim is:

(C1) There is one ultimate moral aim: that outcomes be as good as possible.

C applies to everything. Applied to acts, *C* claims ...

(C2) What each of us ought to do is whatever would make the outcome best.

... Consequentialism covers, not just acts and outcomes, but also desires, dispositions, beliefs, emotions, the colour of our eyes, the climate, and everything else. More exactly, *C* covers everything that could make outcomes better or worse. According to *C*, the best possible climate is the one that would make outcomes best. I shall use ‘motives’ to cover both desires and dispositions. *C* claims

(C5) The best possible motives are those of which it is true that if we have them, the outcome will be best.

As before, ‘possible’ means ‘causally possible’. And there would be many different sets of motives that would be in this sense best: there would be no other possible set of motives of which it would be true that, if we had this set, the outcome would be better. I have described some of the ways in which we

⁹ Feldman is the exception. His theory is much more precise, though quite different in flavour. It shall be discussed after the others.

can change our motives. (C2) implies that we ought to try to cause ourselves to have, or to keep, any of the best possible sets of motives.¹⁰

In ‘Evaluative focal points’, Kagan argues forcefully against forms of indirect consequentialism before introducing his solution:

‘if there is a plausible version of consequentialism, it will evaluate *both* focal points — acts *and* rules — directly. Neither focal point will be elevated to the status of primary evaluative focal point; neither focal point will be evaluated only indirectly.... I have restricted our discussion to two focal points: acts and rules. But as I have already suggested, there are many other evaluative focal points that have been endorsed as *primary* focal points — such as motives, norms, institutions, and decision procedures. As I see it, the most plausible version of consequentialism will indeed be direct with regard to *all* of these. In fact, once we free ourselves from the thought that the evaluative focal points must be at least *prima facie* plausible candidates for the office of *primary* focal point, we realize that absolutely every kind of thing is a potential evaluative focal point (atoms, the weather, sewer systems, suns). So I believe that the most plausible version of consequentialism will be direct with regard to everything.’¹¹

In the same volume, Pettit and Smith introduce their own theory of global consequentialism. Like Kagan they spend most of the article arguing against indirect consequentialism and only a few paragraphs explaining their alternative:

‘Global consequentialism identifies the right *x*, for any *x* in the category of evaluands — be the evaluands acts, motives, rules, or whatever — as the best *x*, where the best *x*, in turn, is that which maximises value.... So, for example, according to global consequentialism, the right act for someone to perform is the act that has greater value than any of the acts that might have been performed instead; the right motive-set for someone to have is the motive-set whose possession has greater value than any of the motive-sets that might have been possessed instead; the right set of rules for someone to have internalised is that set of rules which has greater value than any of the sets of rules that they might have internalised instead, and so on.’¹²

The above accounts have many similarities. In particular, they all advocate that ‘everything’ should be assessed in terms of its consequences. However, there are also a number of differences and ambiguities. By analysing these, we shall be better able to survey the space of possible global consequentialist theories.

¹⁰ Parfit (1984), p. 24–26.

¹¹ Kagan (2000), pp. 150–151.

¹² Pettit and Smith (2000), p. 121.

2.2 Focal points and roles in which to evaluate them

Act-consequentialism allows for the assessment of tokens from the class of acts. To assess a particular act, we examine the outcome that would result from a given agent performing the act at a given time (or over a given period). We can call this the *role* in which the act is assessed. In the case of acts, the appropriate role is obvious and is rarely mentioned explicitly, but what of the other classes that global consequentialism assesses?

When assessing sets of desires, dispositions, beliefs and emotions there seems to be a simple common role. We can look at the outcome that would result from a given agent having the given set of desires/dispositions/beliefs/emotions at the specified time (or over a given period of time). Here the term ‘possessing’ has a subtly different meaning for each of these evaluands, but in each case the correct interpretation seems clear.

However, what are we to say about assessing a given institution or climate? These are not things that an individual can perform or have. In the case of climate, Parfit writes that ‘the best possible climate is the one that would make outcomes best’¹³, but in what role are we to evaluate this climate? Are we to consider it as the climate for Oxford? for England? for the entire world? It is not clear. Moreover, whichever answer we choose, we exclude the other ways of assessing the climate and this seems unnecessarily limiting.

Perhaps we should make the region affected by the climate a parameter in our assessment. We could then assess a climate in terms of the outcome that would result from a given region having the climate over a given period of time. This would work, but it might make us wonder whether there is a single best role in which we are to assess each class of things. The authors above all seem to assume that there is as they provide examples of things being assessed in certain roles without specifying how we are to know which roles to use.

A different approach would be to make our assessments relative to a particular role. Thus we could assess a climate in the role of being had by Oxford in the 18th Century, or being had by Madagascar throughout its existence. This conception of roles is philosophically simple. A role is just a single-place predicate. When given a token from the class under consideration, it forms a proposition. For example, the first role above is that of *being a climate had by Oxford in the 18th Century*. When a climate, such as ‘tropical’ (or something much more specific) is provided, this forms the proposition: *a tropical climate was had by Oxford in 18th Century*. We can then assess the outcome that would result from this being true.

There are three major virtues of the above approach. We have seen two: it avoids the arbitrariness that would otherwise exist in specifying which role we should use to assess tokens from each class and it allows us to use the familiar machinery of

¹³ Parfit (1984), p. 25.

predicates and propositions to make evaluations. The third virtue is that of increased expressiveness.

For example, we could already assess the best set of motives for Jane to have, but now we can also assess the best set of motives for Jane to inculcate within herself, the best set of motives for Rochester to inculcate in Jane, the best set of motives for Jane to inculcate in Rochester, the best set of motives for Jane to aspire to, and the best set of motives for Jane to renounce. Without treating the role as a parameter of the evaluation, we would not be able to express all of these.

This approach is in conflict with a claim made by Pettit and Smith:

‘The intuitive notion of rightness, as applied to acts, is thus just the idea of the act that it is right for someone to perform – as opposed, say, to the act that it is right for someone to forget. In the latter case the evaluand is forgettings, not acts.’¹⁴

However, I don’t think that Pettit and Smith’s alternative can really capture all the same assessments. For example, the best ‘forgetting’ may be of a certain dream, or visual experience and not of an act at all. There would thus be no way to capture the idea of the best *act* to forget. The additional expressivity comes from considering all combinations of focal points with roles in which they could be considered.

While this claim by Pettit and Smith is in conflict with the role approach, it should not be taken as actively denying this method, for it is doubtful that they were explicitly aware that such an approach was possible. Parfit also seems to use one role per focal point, though he too does not consider the alternative. In contrast, Kagan explicitly talks about the possibility of assessing focal points in different roles:

‘But rules ... don’t generate results in and of themselves. Rather, they generate results only when they are ‘embedded’ in some way: they generate results when they are thought about, taught, accepted, disdained, mentioned, mocked, acted upon, flouted, or what have you. So talk of selecting the rules that would have the best results is necessarily shorthand for talk of selecting the rules that would have the best results *if* they were embedded in some specified way.’¹⁵

However, he says this in the context of a rule-consequentialist choosing just one of these ways in which to assess a set of rules. When it comes to the end of the article and the discussion of global consequentialism, he does not re-emphasise this point and it is thus unclear whether he expects the global consequentialist to argue over

¹⁴ Pettit and Smith (2000), p. 133. Note that Pettit and Smith are not explicitly denying the ‘role’ approach that I favour, as it is not clear that they are aware of this possibility.

¹⁵ Kagan (2000), p. 137.

which role (or embedding) is appropriate for which focal point, or to assess focal points relative to any specified role.¹⁶

We thus have two approaches for assessing a focal point. We can do so relative to a fixed role (which depends upon the focal point) and several parameters (such as an agent, a time, a period of time, a region etc), where the list of parameters may itself depend on the focal point. Alternatively, we can assess the focal point relative to an arbitrary role which incorporates the relevant agents, times, or regions in question. We have already seen that the latter is more expressive, and it is also worth noting that it is simpler too, for there is no need to specify a canonical role for each evaluand. Given these advantages, I shall adopt the role approach for the remainder of this article.

2.3 Right or Best?

A key detail in any global consequentialist theory is whether it assesses its focal points in normative terms (right, wrong, should, ought...), or in evaluative terms (good, bad, better, best...). Pettit and Smith make all of their global consequentialist claims in terms of rightness: 'the right act', 'the right motive-set' and 'the right set of rules'.¹⁷ In contrast, Parfit makes global consequentialist claims about 'the best possible climate' and 'the best possible motives'. He makes only evaluative claims concerning non-act focal points, reserving his normative claims for acts: 'what each of us ought to do'.¹⁸ Though it is less clear, Kagan seems to follow Parfit, mentioning 'the best rules' and 'the best motives', but 'the right thing to do'.¹⁹

Let us call those theories that apply normative terms to everything *normative global consequentialism* and those theories that restrict normative terms to acts *evaluative global consequentialism*. Act-consequentialism, as defined in section 1, is silent on the assessment of other evaluands. It could thus be extended into evaluative global consequentialism by the addition of evaluative criteria for all evaluands, or extended into normative global consequentialism by the addition of normative criteria for all evaluands.

In the context of global consequentialism, normative and evaluative terms are not on an equal footing. To my knowledge all global consequentialists accept that if something is right, it is also best. However, movement in the other direction (from the evaluative to the normative) is contentious for focal points other than acts:

¹⁶ One thing that points in favour of the latter is his choice of atoms as a focal point. What could be the canonical role for atoms?

¹⁷ Pettit and Smith (2000), p. 121.

¹⁸ Parfit (1984), pp. 24–26.

¹⁹ Kagan (2000), p. 152.

despite going out of their way to stress the equal status of non-act focal points, Kagan and Parfit still stop short of extending normative terms to cover them.

There are three major concerns for normative global consequentialism. There is the issue of whether it makes sense to apply normative terms to things other than acts, there is the issue of the ownership of the resulting normative claims, and there is the issue of the appropriate type of possibility to comply with the rule that ought implies can. Then there are two concerns for evaluative global consequentialism: there is the issue of whether it involves unnecessarily privileging acts and the issue of whether it is really consequentialist.

The most likely reason for Kagan and Parfit to allow the asymmetry between acts and other evaluands is that they simply do not believe that normative moral terms *can* be applied to things other than acts. This would accord with a common view, that acts are the exclusive bearers of (moral) rightness or (moral) obligatoriness. Many people just believe that it doesn't make sense to talk of the (morally) right set of beliefs to have, or the (morally) right motives. Instead they would see it as some kind of category mistake. They might accept claims of the right beliefs or the right motives if we interpreted 'right' as meaning something like 'correct' or 'rationally required', but not in its moral sense.

However, there is at least some room for disagreement. After all, Pettit and Smith repeatedly apply (moral) rightness to things other than acts and clearly don't think that their claims are nonsensical. It is possible that in their quest for a universal consequentialist theory, they simply confused themselves into committing a category mistake, or an equivocation regarding the sense of the term 'right', but this seems less likely when we consider the way in which normative terms are used in everyday conversation. For example, consider the following 'ought' claims:²⁰

- (1) You ought to be there on Sunday.
- (2) You ought be at work by 9am.
- (3) You ought to be more considerate.
- (4) We ought to be ashamed of ourselves.
- (5) There ought to be a law to deal with this.
- (6) This is the way things ought to be.
- (7) She ought to know her times tables by now

²⁰ I focus on 'ought' for the next few pages because the literature on the related aspects of normativity is all couched in terms of 'ought': 'ought to be' versus 'ought to do'; ownership of oughts; ought implies can. I believe that it could also be carried out in terms of 'should' or 'right'.

(8) You ought to believe in God

(9) Mothers ought to have equal rights

These are quite typical uses of the word ‘ought’ and are all frequently uttered. They are all grammatically correct and they all have fairly clear meanings. It is notable that many of them are in the form ‘ought to be’. This recalls a distinction in normative thinking (particularly in deontic logic) between ‘ought to do’ claims and ‘ought to be’ claims. It is said that these are the two key types of normative claims and that the ‘ought to do’ claims are unique in attaching to a specific agent (the agent whose performance of the act is in question).

One might try to fashion a form of global consequentialism around this distinction, using ‘ought to do’ for acts and ‘ought to be’ for other focal points. For the sake of consistency, it might even be better to forsake ‘ought to do’ and describe acts too using ‘ought to be’. However, I do not think that we are forced into such a move, for I doubt that this distinction really captures what is going on here. Of the above claims, only (5) and (6) seem to be typical of this account of ‘ought to be’. On claims (1)–(3) there is a clearly specified agent who appears to ‘own’ these obligations, despite their being phrased as ‘ought to be’. Claim (4) is similar, though it applies to more than one agent. Defenders of the distinction might argue that these claims are actually disguised ‘ought to do’ claims, which are merely expressed as ‘ought to be’ claims. However, I do not find this convincing, as they do not refer to specific actions and, especially in claims (3) and (4), I do not see how they could be reworded to do so. Finally, in claims (7)–(9) the phrase ‘ought to be’ does not even occur.

Instead, I offer the following explanation of what is going on. To produce a grammatical sentence, we must follow the term ‘ought to’ with a verb.²¹ Most verbs describe actions and are known as ‘action verbs’. Some do not, and these are called ‘state verbs’. An especially versatile state verb is ‘be’, but there are others too. Examples include: ‘have’, ‘possess’, ‘believe’, ‘know’, ‘cope’, ‘flourish’, ‘fail’, ‘succeed’, ‘require’, ‘fit’, ‘resemble’, ‘love’, ‘understand’, ‘need’, ‘want’, ‘prefer’, ‘remember’, ‘forget’ and ‘own’. Many of these are used in everyday ought claims.

We might thus explain the assumption that there are only ‘ought to do’ claims by the mere fact that most verbs are action verbs. We might also explain the reliance on ‘ought to be’ as the only alternative in terms of ‘be’ being the most versatile and prevalent of the state verbs. Indeed, with the heavy pedagogical focus on action verbs (through the common refrain of ‘verbs are *doing* words’), it is understandable that many English speakers are only dimly aware that there *are* state verbs.

If we therefore grant the existence of obligations involving state verbs like ‘have’ and ‘possess’, we can make normative global consequentialist claims such as ‘Oxford

²¹ Note that I am not demanding that our ought claims *be* grammatical (we shall later see that John Broome has good reasons for rejecting this), but that the underlying grammar helps explain why philosophers have focused attention on particular groups of ought claims.

ought to have had a temperate climate in the 18th Century' and 'Alice ought to possess such and such a motive-set' or, alternatively, 'Alice ought to be motivated by such and such a motive-set'.

We might wonder how all this talk of 'ought' connects with 'right'. It is typically considered that they are intertranslatable and that it is a matter of convenience as to which term is given prominence in an ethical theory. The standard consequentialist connection between the two is that (excluding ties) you ought to do something if and only if it is right. In the case of a normative tie between x and y , they are both right, but neither is obligatory. It seems natural to extend this connection to things other than actions. For example, (excluding ties) Oxford ought to have had a temperate climate in the 18th Century if and only if temperate is the right climate for Oxford to have had in the 18th Century.

However, this is one catch: suppose that of all the people Maxwell could kill, the consequences would be least bad if he killed Helen. Then, we would say that the best person for Maxwell to kill is Helen and, on theories like Pettit and Smith's, that the right person for Maxwell to kill is Helen. However, it doesn't follow that Maxwell ought to kill Helen. The problem arises because it is better for Maxwell not to kill anyone at all. This cannot happen with acts, because on consequentialist accounts of acts we must always perform some act, even if it is sitting still for a minute. We can resolve the problem by pointing out that

x is the right y to be in role r

doesn't imply that

It is right for x to be in role r

The connection between them is that:

It is right for x to be in role r iff x is the right y to be in role r and it isn't better to have no y in r than to have x in r

Thus, the correct consequentialist connection between rightness and obligatoriness, is that (barring ties) x ought to be in role r if and only if it is right for x to be in role r . Through this translation, the above defence of normative assessment of things other than acts could be seen as applying to 'right' claims as well as 'ought' claims.

Let us now return to the question that has arisen concerning the ownership of ought claims. When we say 'the right thing for Maria to do is to donate the money' or 'Maria ought to donate the money', we think of these claims as being 'owned' by Maria: they exert a normative pull upon her specifically. Maria is the bearer of the obligation to donate the money. Perhaps others should try to make her do so, perhaps not. That depends on their obligations and not solely upon her own. In contrast, the claims that 'the right climate for Oxford to have in the 23rd Century is temperate', or that 'Oxford ought to have a temperate climate in the 23rd Century' don't seem to be directly owned by anyone. Perhaps we can all play our part in making sure that Oxford maintains a temperate climate into the 23rd Century, but the

obligation does not apply to an individual and it is not clear that it applies to any specific group of people, unless that includes absolutely everyone.

For some claims it is ambiguous as to whether they are owned by anyone. For example, if Pettit and Smith claim that ‘the right set of beliefs for Mary to hold is such-and-such’, does this express an obligation that is owned by Mary? Perhaps its ownership is shared by others who play a significant role in shaping her beliefs, such as her teachers and parents. They might be able to control whether she believes that 12 times 12 is 144 much better than she herself can.

We have seen that the distinction between ‘ought to do’ and ‘ought to be’ is thought to hold the answer here, with all ‘ought to do’ claims expressing ownership and all ‘ought to be’ claims being unowned. However, we have also seen that this account does not seem to adequately deal with all common ought sentences. John Broome has a different approach.²² He denies the standard division and holds that we can separately specify both the obligation and its owner. One of his examples concerns Julie, who ought to have more work to do. It makes perfect sense to consider that Julie owns this obligation and it also makes sense to think that it might instead be her boss’s responsibility and thus that her boss might own this obligation. Broome’s system allows us to express each of these possibilities explicitly.

Broome’s system is also of interest to the global consequentialist for its focus on normatively assessing arbitrary propositions. These propositions may involve actions, such as ‘Maria donates the money’, but they could also express thoughts that don’t involve actions, such as ‘Julie has more work’. While Broome’s system is not committed to consequentialism, it *is* global. One might think of it as a theory of *global normativity*, and we can then see Pettit and Smith’s theory as lying within the intersection of global consequentialism and global normativity.

A potential criticism of Broome’s theory is that it is ungrammatical. His ought claims combine the owner of the claim with the associated proposition as follows:

S ought that *F*

Where *S* is the owner and *F* is the proposition. For instance:

Julie ought that Julie has more work.

The boss ought that Julie has more work.

Maria ought that Maria donates the money.

In his paper, Broome provides a convincing account of why we cannot say such things unambiguously within the grammar of ‘ought’ and yet that we still have such concepts. He sees his expressions of the form ‘*S* ought that *F*’ as a grammatical innovation, that extends the grammar of ‘ought’ and overcomes a merely quirk of the

²² Broome (unpublished manuscript).

English language which was preventing us from verbalising such thoughts previously. Moreover, the focus on arbitrary propositions rather than actions allows for obligations to conditional statements, conjunctions and disjunctions.

There are thus several options for how normative global consequentialists can account for the ownership of oughts. They can let the grammar define the ownership, they can have all oughts be unowned, they can have all non-act oughts be unowned, or they can specify which oughts are owned and by whom. This last route allows more flexibility, but requires an explanation of how oughts come to be attached to specific people in terms of consequences. We shall examine such a theory in section 4.2.

The final issue for the normative global consequentialists is that of possibility. If we accept that 'ought' implies 'can', then there must be some sense in which the obligations of global consequentialism can be met. Since Pettit and Smith do not explicitly account for such possibility in their formula, it may well produce obligations that are impossible to fulfil, such as the right act being to fly away, when the agent cannot fly. The obvious way to remedy this is by imposing some appropriate form of possibility upon the evaluands. For example, Parfit talks of the best possible motives and climates, and specifies this as causal possibility. Pettit and Smith could follow suit by inserting the word 'possible' after the word 'best', as in: 'the right x is the best *possible* x '. A full account would then require an analysis of the appropriate type of possibility, and this would presumably involve the owner of the obligation (if any). However, since there is significant debate about the nature of such possibility, let us just say that we can use whichever type of possibility it is that ends up being involved in ought implies can.

Now we come to the concerns for evaluative global consequentialism. The first of these is that it privileges acts, by singling them out for normative treatment. This criticism holds only if it is possible to normatively assess things other than acts. If this is false (as most evaluative global consequentialists presumably believe), then they are not really privileging acts but are merely assessing each thing in all ways that it can be morally assessed.

The second concern is that evaluative global consequentialism is not really consequentialist, for it is quite compatible to say that ' x is the best set of motives for Jane to have', ' x is a better set of motives for Jane to have than y is' and yet ' y is the right set of motives for Jane to have'. This mirrors the non-consequentialist's claims that ' z is the best act', but ' z is the wrong act'. Perhaps the evaluative global consequentialist is agreeing that Jane's possession of motives x would lead to more good, but thinks that motives y are more godly, and thus right. Again, this criticism only really holds if it is possible to have normative assessment of things other than acts. If not, then the non-consequentialist possibility is no longer assertable, and the problem vanishes.

In conclusion, it seems to be an open question as to whether we can normatively evaluate things other than acts. If we can't, then evaluative global consequentialism appears to be able to provide a satisfactory answer to the rough question, and normative global consequentialism is impossible. If we can, then we should use a

version of normative global consequentialism. In this case, we must decide on an appropriate theory of possibility to take account of how ought implies can and we must also decide how the resulting obligations are owned, if at all.

At this point it is worth mentioning an intriguing possibility. Some recent authors have argued that consequentialism is best considered as a theory that uses only evaluative terms and eschews normative terms altogether. Frances Howard-Snyder and Alastair Norcross have named this *scalar consequentialism*.²³ Taking a similar approach to global consequentialism would produce scalar global consequentialism: a theory which evaluates everything but applies no normative terms at all. It would thus treat acts and non-acts exactly alike without needing to apply normative terms to the non-acts. We shall not explore scalar global consequentialism further here, but merely note the possibility.

2.4 Formalizing role-based global consequentialism

We can now take the above considerations into account and formulate principles for a role-based theory of global consequentialism. We shall start with a principle of betterness:

Global consequentialist criterion of betterness within a role

x is better than y in role r iff the outcome that would occur were x in role r , is better than the outcome that would occur were y in role r

Stated clearly like this, we can see a problem for such a principle of betterness which has not been remarked upon by other authors. Whereas acts are often defined in such a way that two alternative acts cannot both be performed, we have not yet blocked this possibility here. Such a situation can lead the above principle to make unintuitive judgments. For example, suppose we are considering two events — *the holocaust* and *the scratching of my finger* — within the role of *happening*. Now suppose that I don't actually scratch my finger. In this case, the outcome were the holocaust to happen is just the actual historical outcome, while the outcome were I to scratch my finger is an outcome that is slightly worse than the actual outcome (for it includes both the holocaust and the scratching of my finger). Thus, if applied to this case, the above principle would rate the scratching of my finger as worse than the holocaust.

The problem stems from the two things not being proper alternatives for the role: one is in the designated role whether or not the other is. To get around this, we could make any of the following stipulations:

- (1) The principle is only to be applied when it is impossible for x and y to both be in r

²³ Howard-Snyder and Norcross (1993).

(2) The principle is only to be applied when, were x to be in role r , y would not be and were y to be in role r , x would not be.

(3) Roles must be specified such that no two things can simultaneously fill them.

Stipulation (1) restricts betterness judgments to mutually exclusive things, and thus avoids the problem. Stipulation (2) also avoids the problem, but does so without sacrificing as much expressiveness. For example, it lets us compare two events (both of which could happen) so long as were one to happen, the other wouldn't. Both (1) and (2) are ways of saying that x and y must be *alternatives*. We could say that (1) restricts x and y to being *hard alternatives*, while (2) restricts x and y to being *soft alternatives*. Stipulation (3) is quite similar to (1), but puts the burden on the specification of roles. For example, it wouldn't allow the role of *being a motive possessed by Susan*, as there are outcomes in which Susan possesses more than one motive, but it would allow *being the motive-set possessed by Susan*. Since (2) is the most expressive, it is my preferred option.

It is also worth noting that the global consequentialist criterion of betterness involves a counterfactual and thus requires some explanation of the type of possibility involved. My approach is to understand this as a very liberal type of possibility, potentially even beyond physical possibility. For example, one could ask 'would it be better for Alex to have the ability to fly or the ability to become invisible?' Even though both abilities might be physically impossible, we can still make sense of this question relative to a shared background of how such supernatural abilities are supposed to work. By using such a liberal type of possibility, we maximise the expressiveness of the theory.

We can now add a principle to say when something is the best within a role:

Global consequentialist criterion of best within a role

x is the best e in role r iff the outcome that would occur were x in role r , is better than the outcomes that would occur were any other e in role r

The ' e ' in the above principle stands for any category of evaluand (such as *acts* or *institutions*), whereas the ' x ' is a member of this category (such as *sitting down* or *the House of Lords*). For example, one might use the criterion to ask whether the House of Lords is the best institution for legislative review in England. While we defined betterness without reference to the evaluand, it is needed here to provide a class of things with which to compare x .

This principle is very similar to the formulations of Parfit and Kagan, although it explicitly allows focal points to be assessed with respect to different roles. Together, the two principles we have formulated provide a basis for making evaluations, but do not allow us to make any normative claims (not even for acts). If we want a theory of evaluative global consequentialism, we can add the act-consequentialist criterion of the rightness of acts (from section 1). This would match Parfit and

Kagan's approaches, and would provide a specific theory of role-based evaluative global consequentialism.

To form a theory of normative global consequentialism, we could add the following principle instead of the act-consequentialist criterion of rightness of acts:

Global consequentialist criterion of rightness within a role

x is the right e in role r iff x can be in role r and for each other e that can be in role r , x is better than it in r .

This gives a version of normative global consequentialism that is similar to Pettit and Smith's, but also allows for assessment with respect to arbitrary roles.

It is important to note that in this principle, the conception of possibility must be much more limited than for those above. While we can sensibly talk about 'the best superpower for someone to have' even though any superpower would be a physical impossibility, it would be another thing altogether to talk about 'the right superpower for someone to have', for this would presuppose that they *can* have a superpower in a sense relevant to the injunction that ought implies can.

We might therefore imagine that we are dealing with an individualistic and physical sense of can, such as 's can do x if and only were s to desire to do x , s would do x '. However, there are sensible sounding rightness claims that go beyond this level of possibility. For example, that 'the current climate is the right climate for the world to have in 2100'. It may be that there is no change to an individual's desires which would reverse climate change, but a change to multiple individuals' desires would suffice.

In its current form, this principle is somewhat ambiguous between the individual conception of possibility and the collective conception: if we used it to produce a claim about the right act for someone to perform then we would naturally use a conception of possibility relative to that person, while for claims about the future climate we would naturally assume a collective conception. If we wanted, we could remove this ambiguity by explicitly mentioning the owner of the resulting obligation. This can be achieved by switching to Broome's terminology for owned ought claims, and adding a third clause to the right hand side.

Global consequentialist criterion of role-based obligations

s ought that x be in role r iff x can be in role r ; for any other thing that can be in role r , x is better than it in role r ; and it isn't better to have nothing in role r .

With this new version of the principle the conception of possibility is tied to s (the owner of the obligation).

3. Proposition-based global consequentialism

3.1 Feldman's global consequentialism

So far we have looked at the global consequentialist theories of Parfit, Kagan, Pettit and Smith. We have seen the details that need to be filled in and several ways in which this could be done. However, there is a theory of global consequentialism that is already complete. In 1975 Fred Feldman wrote a paper called 'World Utilitarianism' in which he advocated a new form of consequentialism that used possible lives an agent could have led as its focal point. His aim was to avoid technical problems for the contemporary formulations of act-consequentialism and he used his theory to assess acts in a novel, and consistent, manner. In his 1986 book, *Doing the best we can*, Feldman kept the basic structure of world consequentialism, but extended it to assess arbitrary propositions rather than just acts. Then, in his 1993 paper 'On the consistency of act- and motive-utilitarianism: a reply to Robert Adams', he explicitly showed how his theory could be used as a form of global consequentialism (which he called 'wide consequentialism').

World consequentialism starts with the following theory of possibility:

'[A]t every moment of moral choice, there are many possible worlds *accessible* to the moral agent. Consider some infant at the very moment when she begins to exist. There are various possible ways in which she might live out her life. For each of these total histories, there is a possible world — the world that would exist if she were to live out her life in that way. Call these worlds the "original stock" for the person. As time goes by, the person filters out world after world. By acting in this way rather than that, she bypasses, or rules out, many possible worlds. Once bypassed, a given world is never again accessible to a person. If we select a person, *s*, and a time, *t*, at which *s* exists, there will be a not-yet-ruled-out subset of *s*'s original stock. Each member of this subset is a possible world that is, in an important sense, still possible for *s*. I say that each such world is "accessible" to *s* at *t*.'²⁴

To this Feldman added a definition of a 'best':

A world *w* is a *best* for *s* at *t* iff *w* is accessible to *s* at *t* and no better world is accessible to *s* at *t*

Now we can define the fundamental principle of world consequentialism:²⁵

²⁴ Feldman (1993), p. 208.

²⁵ We use here the principle as outlined in Feldman (1993). It seems likely that Feldman still prefers his more careful statement of the principle from *Doing the best we can* which traded elegance in formulation for the ability to deal successfully with cases in which there an infinite succession of accessible better worlds and no best ones. The reader should feel free to substitute in that version if so desired — nothing will turn on it here.

World-consequentialist criterion for obligatory propositions

As of t , s morally ought to see to the occurrence of p iff p is true in all the bests for s at t

In this formulation, p is an arbitrary proposition. For example, p might be the proposition expressed by 'Maria donates the money' or that expressed by 'Oxford has a temperate climate in the 23rd Century'. Thus if it is true that Maria donates the money in all the bests accessible to Maria at time t , then as of t , Maria ought to see to the occurrence of Maria donating the money (i.e. Maria ought to donate the money). If it is true that Oxford has a temperate climate in the 23rd Century in all the bests accessible to Maria on 1/1/2008, then Maria ought to see to the occurrence of Oxford having a temperate climate in the 23rd Century as of 1/1/2008.

There are a number of interesting aspects of Feldman's theory. For one, it is a tensed theory of moral obligation. An act being performed at time t might be obligatory as of one time and impermissible as of another. This may seem odd, but it can be a desirable feature. An example shows why.

The followup

Jacinta is suffering from a deadly medical condition and has come to see her doctor, Valerie. There are two ways to treat the condition. One is to give a dose of drug A now and another dose of drug A in a followup procedure in a week's time. The other is to give a dose of drug B now and another dose of drug B in a followup procedure in a week's time. Giving the two doses of drug A is slightly better than the two doses of drug B because it has fewer side effects. Mixing the drugs, or not giving the drugs at all will lead to Jacinta's death.

As of Jacinta's arrival at the hospital, Valerie has it open to her to provide any combination of the two drugs (or no drug at all). In the best possibility, Valerie provides drug A both times. Thus on Feldman's view, and according to common sense, she ought to administer drug A now and she ought to administer drug A at the followup appointment. But suppose that Valerie actually administers drug B in the first case. After this, it appears false that she ought to administer drug A in the followup. We can overcome this problem with a theory of tensed rightness such as Feldman's. At the time of the first meeting, Valerie ought to administer drug A at that meeting and ought to administer drug A in the followup. Since she actually administers drug B at the first meeting, at all times after that, she administers drug B in the best worlds open to her and thus *as of that time*, she ought to administer drug B in the followup.

The tensed nature of Feldman's theory can thus be seen as a virtue. Indeed a major reason for the lack of tensed rightness in other theories (such as traditional act-consequentialism) is that in most cases the judgment about rightness is made at the same time that the act is performed. However, it is clear that we can ask about situations where this is not the case and a tensed theory can resolve the ensuing issues successfully.

There is a strange quirk of Feldman's theory which is related to its tensed account of obligation. Any proposition becomes obligatory once it can no longer be made false, for it is then guaranteed to be true in all accessible worlds and thus all bests. For example, suppose Michael murders someone at time t . As of any time after t , the proposition expressed by 'Michael murdered someone at t ' will be true in all worlds accessible to Michael and as of any time after t , Michael *ought* to see to the occurrence of Michael murdering someone at t . In short, the issue is that anything that can't be avoided by s at t is obligatory to s at t .

This is not obviously false, and Feldman accepts it as an odd consequence of his theory. However, we may think that just as 'ought' implies 'can', so 'ought' implies 'can fail to'. Feldman caters to this possibility by showing how a small modification can be made to his theory to remove these troubling obligations and adhere to this stronger principle. We can define the concept of a proposition being within someone's power as follows:

p is within s 's power at t iff p is true in at least one world accessible to s at t , and
 p is false in at least one world accessible to s at t .

We can then add a clause requiring that a proposition must be within the agent's power:

World-consequentialist criterion for obligatory propositions (revised)

As of t , s morally ought to see to the occurrence of p iff p is true in all the bests
for s at t and p is within
 s 's power at t

This revised formulation also deals with another strange case. First, let us recall the case of Julie who ought to have more work. Feldman's theory can tell us that if Julie has more work in all the bests accessible to her, then she ought to see to the occurrence of Julie having more work. It can also tell us that if Julie has more work in all the bests accessible to *Julie's boss* then Julie's boss ought to see to the occurrence of Julie having more work. Thus, the theory express ownership of oughts in the flexible way discussed by Broome: the owner need not feature in the proposition.

However, in its original form world consequentialism gives too many owned oughts. For suppose that Julie's boss actually will increase Julie's work and consider the obligations of a third party, George, who lives in a different country and cannot influence how much work Julie will have. Since Julie gets more work in every best accessible to George, on the original formulation this implies that George ought to see to the occurrence of Julie having more work. On the revised formulation, however, only those who have accessible worlds where Julie gets more work and accessible worlds where she doesn't can have such an obligation. It thus seems to deal with ownership of obligations in just the right way. Since the revised formulation deals with both the counterintuitive oughts regarding past events and the counterintuitive oughts regarding other people's affairs (and indeed counterintuitive oughts involving anything out of one's power), we shall accept the revision from here on.

World consequentialism is a form of global consequentialism. It assesses acts, motives, climates, eye-colours (and everything else) in terms of the goodness of the outcomes in which they occur. In Feldman's words:

'[World consequentialism] may be taken as a global theory of normative evaluation. It tells us what actions we should perform, what motives we should have, what traits of character we should develop, etc.'²⁶

World consequentialism is also the most developed variety of global consequentialism. It has a consistent set of answers to each of the concerns that were discussed earlier in this chapter. Focal points (acts, motives etc.) are each assessed in the context of a given role, and this is achieved via propositions that combine the two. For example we assess a given set of motives in the role of being had by Jane during May, using the proposition that *Jane has the motive-set during May*. Since propositions are very fine-grained, the theory can express a great deal. The theory also benefits from the fact that propositions are so familiar and well-studied. For example, we do not need to provide a detailed theory of action to assess acts, a detailed theory of motivation to assess motives and so forth.

World consequentialism is a variety of normative global consequentialism. It resolves the questions of ownership of obligations in a flexible way, and it provides an account of the possibility involved in 'ought' implies 'can' (or at least a sketch of the appropriate possibility). Interestingly, it is a deeply normative theory in that it does not lend itself to a single obvious way of merely *evaluating* propositions.²⁷

Given the long history of world consequentialism and its degree of sophistication, it is interesting that it has been ignored in recent discussion of global consequentialism. One reason for this might be that its approach is quite different to the three varieties discussed earlier. Whereas Parfit, Kagan, Pettit and Smith sought to evaluate an act (or climate, or set of motives) in comparison to other possible acts (or climates, or sets of motives), Feldman compares a world to other possible worlds and then looks at the acts (or climates, or sets of motives) that occur within these worlds.

This difference in strategies leads to an important difference in the type of theory. Parfit, Kagan, Pettit and Smith all have theories that are compatible with both actualism and possibilism. They could define the best motives as those which will actually have the best results (focusing on the outcome where everyone acts as they actually will), or they could define the best motives as those which are required in order to reach the best possible results (focusing on the outcome where the agent in

²⁶ Feldman (1993), p. 212.

²⁷ One could perhaps evaluate a proposition in terms of the goodness of the best world in which it occurs, but this would have to be a concept of the goodness of the proposition *relative to an agent s* and we typically think of the goodness of something as being agent-independent.

question chooses the best life available to her). In contrast, Feldman's theory is possibilist by its very nature, for its prescriptions are based solely around what happens in the best worlds accessible to the agent.

There are a number of ways in which the actualist conception of obligation better fits our intuitions. For example, when we advise our friends and tell them what they ought to do, we do better if we take into account the actual consequences of their acting rather than assuming that they will produce the optimal outcome associated with their action. More tellingly, suppose you are considering a choice of your own where option *A* could potentially lead to a tiny amount of good, but only if you followed through and were strong enough of will to avoid many pitfalls. If you were weak of will even once, it would lead to a great calamity. In contrast option *B* is safe and boring, leading to no additional good or harm. Let us also assume that you are quite weak of will and that, while you *could* avoid all the temptations on the first course of action, you almost certainly won't. It seems that it is best to settle for option *B*. However, possibilist theories like Feldman's say that you ought to choose option *A* because the best accessible outcomes involve it. This obligation sounds acceptable in theoretical terms, but it would cause great problems if it were action guiding.

This is not a knock down objection to possibilism, but shows how it clashes with certain consequentialist intuitions and thus why many consequentialists might seek an actualist theory of global consequentialism. Let us therefore construct an actualist theory that combines some of the virtues of world consequentialism with an underlying intuition more in keeping with the approaches of Parfit, Kagan, Pettit and Smith.

3.2 An actualist account

Following Feldman, we shall use propositions as the objects of our assessments. However, we shall have to change the method by which we assess propositions. When actualist act-consequentialism evaluates acts, it does so in terms of *what would happen* if the act were performed: not what *might* happen, or what would *ideally* happen, but what would *actually* happen. However, when world consequentialism assesses a proposition *p*, it does so by comparing the best accessible world in which *p* is true and the best accessible world in which *p* is false. To create an actualist theory, we must harmonise the assessment of propositions with that of acts. The obvious way to do so is to compare what would actually happen if *p* were true with what would actually happen if *p* were false. More generally, we could say:

Global consequentialist criterion of betterness for propositions

p is better than *q* iff the outcome that would occur if *p* were true is better than the outcome that would occur if *q* were true.

Thus, following actualist act-consequentialism, we would conclude that *Maria donates the money* is better than *Maria spends the money on chocolate* if and only if the outcome that would result from Maria donating the money is better than that which would result from Maria spending the money on chocolate.

This principle has much in common with our earlier principle of betterness relative to a role, and like that principle, it must be restricted to cases where p and q are alternatives to each other. As before, we can either use the more familiar concept of hard alternatives (which cannot both be true) or the more expressive concept of soft alternatives (were one true, the other would be false).

Note that this criterion for the betterness of propositions is even more expressive than the earlier criterion for betterness within roles. This is because the two propositions can refer to different roles. For example, we could compare the proposition that Clare cares deeply for her child with the proposition that Clare benefits a stranger at the expense of her child.²⁸

From here, the natural move is to define obligations in terms of the ‘best possible’. However, this was much easier when we were looking for the best member of a given class. When it comes to propositions, we can only move from better to best in the context of a set of alternatives. Rather than attempt to specify the appropriate alternative set, let us begin by looking merely at an agent’s obligations relative to a given set of alternatives.

Global consequentialist criterion of ought-rather-than for propositions

As of t , s ought that p rather than $q_1 \dots q_n$ iff $p, q_1 \dots q_n$ are each within s ’s power at t , and p is better than each of $q_1 \dots q_n$

This principle bears many similarities to world consequentialism: it uses a tensed ought, it separates the owner of the obligation from the contents of the obligation and it specifies ‘can’ in terms of the worlds accessible to the agent (via the former definition of ‘within one’s power’). However, unlike world consequentialism, this is an actualist theory. For example, suppose that Brian is deciding whether to stay home or to go to the pub. He hopes to work if he stays home, but will instead slump in front of the television. The outcome of staying home and working is the best, followed by going to the pub, followed by staying home and slumping.

This principle claims, ‘Brian ought that *Brian goes to the pub* rather than *Brian stays home*’ because it is within Brian’s power that he goes to the pub and the outcome that would result from him going to the pub is better than the outcome that would result from him staying home (and consequently slumping in front of the television). Interestingly, this principle also claims that ‘Brian ought that *Brian stays home and works* rather than *Brian goes to the pub*’. Contrary to appearances, there is no conflict between these ought claims because they are made relative to different alternatives.

Before we provide the last piece in this global consequentialist theory, let us consider just what is missing. Suppose that we want to know whether Maria ought that Maria donates her money. While our theory does not answer this, it does answer a lot of very similar questions. In particular, for each set of alternatives S , it tells us whether

²⁸ To see why this might be particularly useful, see Parfit (1984), pp. 32–5.

Maria ought that Maria donates her money rather than any alternative in S . Is this enough? In one sense it seems that it is. For suppose that the relevant set of alternatives is some set of actions S^* . Now one of the things that the theory already tells us is whether or not Maria ought that Maria donates the money rather than any alternative in S^* . It just does not tell us this under a special description where the alternative set is not specified. If we knew which were the relevant alternatives, then we would not need to know any more, and arguably this question of relevance is the asker's business and not the moral theory's business.

I thus propose the following analysis of ought claims where no alternatives are specified:

's ought that p ' ambiguously stands for some claim of the form: s ought that p rather than $q_1 \dots q_n$. There is no single correct way to resolve it as it depends on speaker and context.

Thus, on this account we could say both that 'Brian ought to stay home and work' and that 'Brian ought to go to the pub' without any explicit conflict, for the obvious interpretation is that these prescriptions are made from different implicit alternative sets. It is worth noting that this implicit approach to alternatives finds support even outside the traditional actualist/possibilist trouble cases. For example, someone might be in a situation where she promised to meet someone across town, but has just run into an old friend and asks you what she should do. Suppose you reply that she should go to the post office and post a cheque for £10,000 to Oxfam, because this is the best thing she could possibly do now. She may well reply that this would indeed be best, but it wasn't one of the implicit alternatives she was considering. She wanted to know whether she ought keep her promise rather than catching up with the old friend, not whether she ought keep her promise rather than anything else she could possibly do at the time. Understanding ought claims with no specified alternatives as ambiguous between different alternative sets, rather than pointing at the one true alternative set, helps to explain such cases.

Conclusion

Since the classical utilitarians, philosophers have wanted to apply consequentialism to everything. More recently, they have suggested a number of ways in which this could be done. By exploring the differences in their approaches and developing new theories of global consequentialism in response, I have tried to further this goal.

References

- Adams, Robert, 1976. 'Motive utilitarianism', *Journal of Philosophy* 73:467–481.
- Bentham, Jeremy, 1789. *An introduction to the Principles of Morals and Legislation*, (New York: Hafner), 1961.
- Broome, John, 'Ought', unpublished manuscript.

- Feldman, Fred, 1975. 'World utilitarianism', in Keither Lehrer (ed.) *Analysis and Metaphysics*, (Dordrecht, Reidel), 255–71.
- Feldman, Fred, 1986. *Doing the best we can*. (Dordrecht: Reidel).
- Feldman, Fred, 1993. 'On the consistency of act- and motive-utilitarianism: a reply to Robert Adams', *Philosophical Studies* 70:201–212.
- Howard-Snyder, Frances, and Alastair Norcross, 1993. 'Rejecting the right', *Journal of Philosophical Research* 18:109–125.
- Kagan, Shelly, 2000. 'Evaluative focal points', in B. Hooker, E. Mason and D. E. Miller (eds.) *Morality, Rules and Consequences: A Critical Reader*, (Edinburgh: Edinburgh University Press), 134–155.
- Mill, John Stuart, 1843. 'A system of logic: ratiocinative and inductive', in John M. Robson (ed.) *Collected Works (vol VIII)*, (Toronto: University of Toronto Press), 1974, 1–952.
- Parfit, Derek, 1984. *Reasons and persons*, (Oxford: Clarendon Press).
- Pettit, Philip and Michael Smith, 2000. 'Global consequentialism' in B. Hooker, E. Mason and D. E. Miller (eds.) *Morality, Rules and Consequences: A Critical Reader*, (Edinburgh: Edinburgh University Press), 121–133.
- Sidgwick, Henry, 1907. *The methods of ethics*, 7th edition, (London: Macmillan).