

Global consequentialism and the morality and laws of war

Hilary Greaves, Oxford

Draft, 2 March 2016

1. Introduction

Rights-based approaches and consequentialist approaches to ethics are often seen as being diametrically opposed to one another. This is entirely understandable: *simple act-consequentialism* (at least) just is the thesis (roughly) that one has moral reason to perform some action to the extent that performing the action in question tends to lead to the best consequences overall, while (on the other hand) to say that X has a moral right to Y is in part to assert that there are moral reasons for suitably situated agents to provide X with Y *even if doing so foreseeably will not* lead to the best consequences overall. At this level, intuitions in various cases are strongly on the side of the rights-theorist: intuitively, for example, one ought not to execute an innocent scapegoat, even if doing so will foreseeably lead to better consequences overall thanks to its role in pacifying the mob.

On the other hand, we may well, like Narveson (1973, p.69), “see nothing but darkness and confusion” in a view that takes rights as literally primitive – a view, that is, according to which there are simply *facts* both about which persons have rights to which things and what this implies about what one ought to do, which facts are (on this view) not themselves grounded *in any sense* in any kind of reasons, either for ascribing these rather than those rights or for responding to rights in the prescribed way.

Global consequentialism aims, among other things, to find the appropriate middle ground between these two extremes: simple act-consequentialism at one extreme (with no place at all for talk of rights), and primitive rights-theory on the other (with no place at all for rationales for rights-ascriptions or rights-respect). Global consequentialism is also an independently motivated view: In particular, it grows naturally out of familiar concerns that simple act-consequentialism, if construed as the whole of morality, might be self-defeating. It is, in any case, *one* natural account of the foundational morality that might underlie rights-talk, and hence that might have distinctive implications for how disagreements at the level of practical moral and legal rights-theorising are to be resolved.

The purpose of the present chapter is to illustrate how this might work: I will apply global consequentialism, by way of case study, to a recent discussion of the ethics of war. Section 2 reviews the background theory of global consequentialism. Section 3 reviews a recent exchange between Jeff McMahan and Henry Shue over the morality and laws of war. Section 4 is an attempt to work out what a global consequentialist should say about this particular exchange: I will argue that by global consequentialist lights, each of McMahan and Shue is in some important respects clearly right, in others clearly wrong.

2. Global consequentialism

The rights-violation objection. The basic objection to consequentialism based on considerations of *rights*, or (more generally) side-constraints, is familiar. Suppose, for example, that a town is gripped by anger at a spate of recent unsavoury crimes, and that the town sheriff faces the choice between

executing an innocent scapegoat (which would pacify the mob) on the one hand, or allowing the riots to continue (in which case many more people will be killed) on the other. At least in hypothetical cases in which, we stipulate, there will be no knock-on consequences of the sheriff's decision on e.g. general respect for the machinery of justice, simple consequentialism will say that the sheriff ought to execute the innocent scapegoat. But most people find this implausible. The scapegoat, intuition screams, has a *right* to be treated fairly, it is *unjust* to punish the innocent; perhaps with the exception of cases in which the consequences at stake are really extreme, these considerations of rights and justice override considerations of overall good in determining what the sheriff ought to do.

The self-defeatingness objection. Aside from the brute intuition against executing scapegoats for the sake of the greater good, however, simple act-consequentialism is also subject to a quite different kind of objection: that it is *self-defeating*. The latter worry is that trying to *use* act-consequentialism to guide one's decisions on an everyday basis, and/or to formulate laws that are themselves directly consequentialist in content, would itself lead to undesirable consequences, so that the consequentialist *even by her own lights* should want to avoid such 'consequentialist' decision procedures, laws and so on.

There are at least three ways in which this can happen. First, relating to calculation time: trying to work out, in full gory detail or even to a reasonable approximation, the possible or likely consequences of one's proposed action from here to eternity, together with probabilities for all the relevant possibilities, is clearly an unwieldy task, and not one to which the computational power of human brains is always well adapted. For this reason alone, it would lead to bad consequences if we actually tried to perform full consequentialist calculations before, for example, deciding whether or not to save a drowning child. We will get better consequences if we have adequately reliable 'rules of thumb' worked out in advance, and if in the heat of the moment, we simply apply those rules without (much) further reflection.

Second, relating to personal bias: Because full consequentialist calculations are complex and hence (in practice) require a fair amount of estimation, there is a lot of scope for personal bias to skew the results, if we make case-by-case decisions on the basis of attempts at consequentialist calculations. Suppose, for example, that I have promised an elderly and somewhat tiresome aunt that I will call her at 8pm this evening, but then I receive a last-minute invitation to an appealing-sounding party, a prospect that I vastly prefer on self-interested grounds. A consequentialist decision procedure, it seems, would have me try to estimate the amount of harm it might do to my aunt if I break my promise, and the likelihood of that harm materialising, so that I can compare these possible consequences with my estimates of the enjoyment and other benefits the party will bring me. But clearly, I have an incentive to underestimate the former and overestimate the latter, so as to get the calculation to deliver the verdict that I ought to go to the party: therefore, there is a danger here of motivated cognition. I might in general do a better job of choosing actions with the best consequences if my decision procedure is something other than *thinking* about which actions have the best consequences.

Thirdly, relating to alienation: most or all of us are not psychologically well-suited to conducting consequentialist reasoning in our daily lives. Arguably, for instance, we naturally go to visit a close friend or family member in hospital driven by values that are far from impartial: driven, by, for

instance, a direct concern for the welfare of the hospitalised party in particular, an enjoyment of his/her company, a high value placed on our own special relationship to this person. While there might, in the end, be reasons that even the consequentialist should accept for visiting such a person rather than (say) using the time in question to volunteer for charity – perhaps one’s ability to cheer up close friends means that the hospital visit would do more good *even considered impartially* than one could do by serving in a soup kitchen, and/or perhaps one needs to maintain such relationships in order to stay psychologically healthy enough to be productive in one’s charitable works in the long run – the very process of running all our decisions through such an explicit consequentialist algorithm might lead to a sort of “alienation” between one’s affections and one’s deliberative self – one that in turn leads to a sense of loneliness and emptiness, is destructive to valuable relationships, and leaves the world overall in a worse state than it would have been in had some not-explicitly-consequentialist decision procedure been followed (Railton (1984)).

This discussion has been carried out in terms of the everyday decisions of private individuals, but it is if anything even clearer that similar considerations apply to the question of (e.g.) what the content of our laws ought to be. There are all sorts of reasons for thinking that the consequences of having laws that simply read “do whatever would have the best overall consequences” would itself be disastrous (relating, for instance, to epistemic and motivational limitations, both on the part of the courts and on the part of those who would be subject to the law in question). Even the consequentialist, then, should not desire laws that are themselves consequentialist *in content*. In *that* sense, “consequentialism about the law” would be self-defeating.

Global consequentialism: a first pass. This much seems clearly correct: making all of one’s everyday decisions on the basis of explicit and case-by-case consequentialist calculations, and/or having laws that themselves simply restate consequentialism, would generally have seriously suboptimal consequences, and so even by the consequentialist’s own lights would be undesirable.

It would be too quick to conclude from this, however, that consequentialism itself must simply be rejected wholesale, as internally inconsistent, and therefore a false and misguided ethical theory. The warranted conclusion is more limited: it is simply that the consequentialist, by her own lights, should prefer decision procedures, laws, and so forth that are not themselves simply-mindedly consequentialist in content. But of course, this does not render considerations of consequences altogether irrelevant (that is, consequentialism does not, *pace* Williams (1973, p.134), hereby “usher itself from the scene”). The consequentialist should prefer the enactment of whatever set of laws has the property that enacting *that* set of laws would lead to better consequences than enacting any alternative set of laws; she should prefer the adoption of whatever everyday decision procedure has the property that adopting that decision procedure has better consequences than adopting any alternative decision procedure; and so on.

Further: We now see how a framework of rights, in particular, might in one clear sense be reconciled with consequentialism. Granted that the laws, everyday decision procedures and so forth whose adoption leads to the best consequences are unlikely (for the above reasons) to be consequentialist in content, in particular, there is no reason to think that their content could not involve a rights-framework. Hence, there is no reason of principle to doubt that global consequentialism might amount to the middle ground, between simple act-consequentialism and primitive rights-theory, whose desirability I suggested at the start of this chapter. In this approach, rights might find a place

at several levels of morality (so that the initial intuitions against simpleminded consequentialism are accommodated), but the *foundations* for all levels are ultimately consequentialist ones (so that “darkness and confusion” is avoided).

The structure of normative theories. To formulate global consequentialism more generally and (slightly) more precisely¹, it is useful (roughly following Kagan (1992)) to distinguish between *evaluative focal points*, *roles* and *normative factors* in the formulation of a normative theory.

- A theory's *evaluative focal point* is the thing whose normative status the theory seeks to assess. We have already considered acts, laws and everyday decision procedures; one might similarly evaluate character traits, motivations, and so forth. Simple act-consequentialism, for example, is a normative theory that seeks to assess only one type of focal point (acts), but other normative theories may be concerned with more or different evaluands.
- The *normative factors* are the factors that the theory considers relevant to the normative status of the evaluative focal point. What is distinctive of *consequentialist* theories is that they, and they alone, admit only the goodness of consequences as a relevant normative factor. In contrast, a Kantian theory might admit, for instance, universalizability (in some more or less precisely specified sense), in addition to or instead of considerations of consequences.
- The need for specification of a *role* arises because there is, in general, more than one thing one might ‘do with’ the item being evaluated, and in general the evaluation will depend not only on the item, but also on what is done with it. For example, the best act to *try to perform* may not be the same as the best act to *perform*, since tryings can have other consequences besides performance. Similarly, the best set of laws to advocate for may not be the same as the best set of laws to enact; and the best set of moral rules to guide one’s own actions by may not be the same as the best set of moral rules to use as a basis for ascriptions of praise and blame.

Using this scheme, we can classify (in particular) act-, rule- and global consequentialism, and see how they differ. Since these three theories are all forms of consequentialism, they agree about the ‘normative factors’; but they disagree along the other two dimensions. Act- and rule-consequentialism each choose some *particular* evaluative focal point (respectively, acts and ‘rules’), and evaluate their chosen focal point relative to a particular role. The relevant role for act-consequentialism is of course the act’s being performed; rule-consequentialists disagree among themselves about the choice of role, but any particular rule-consequentialist theory takes on some particular commitment on this matter (for instance, the form of rule-consequentialism advocated by Hooker (2000, chapter 3) takes the relevant role to be: being ‘accepted’ by a majority of members of the moral community, with some discussion of exactly what is involved in ‘acceptance’).² In contrast, global consequentialism declines to single out either any particular evaluand(s) or any particular

¹ Unlike act- and (to a lesser extent) rule-consequentialism, there is no canonical, widely accepted formulation of global consequentialism. Related ideas have been discussed in: Adams (1976), Parfit (1984, chapter 1), Feldman (1986), Railton (1988), Crisp (1992), Feldman (1993), Kagan (2000), Pettit and Smith (2000), Ord (2008).

² Rule-consequentialism then adds a further, *non-consequentialist* criterion for the evaluation of acts, viz. that an act is right iff it is in conformity with the selected rules.

role(s) as primitive: as the name suggests, it amounts to consequentialist evaluation of *any* evaluative focal point, relative to *any* role.

Global consequentialism: official formulation. Global consequentialism can then be formulated, slightly more precisely, as follows:

(GC) Global consequentialism: *For all evaluative focal points F and roles R, the right F in role R is the F such that having that F in role R has better consequences than having any alternative F (or no F at all) in role R.*

(GC) entails, for instance, that the right *act to perform* is the act whose performance has better consequences than the performance of any other available act (or no act); and that the right *act to try to perform* is the act such that trying to perform that act has better consequences than trying to perform any available act (or no act); and so forth.

So far so good, perhaps; it is a slightly delicate matter, though, what exactly this has to do with such questions as that of whether or not there is a moral right to life. For sure, it follows from the above that according to global consequentialism, the 'right' set of items to *say* that there are moral rights to is that set of items such that saying there are moral rights to those items has better consequences than saying there are moral rights to any alternative set of items (or to no items). That is, we can see that global consequentialism will take a consequentialist line about what it is 'right' to say. But for all that has been said so far, this could simply be the familiar line taken by act-consequentialism, viz. that it is morally right to say whatever will have the best consequences, regardless of truth or falsity. And, returning to the dilemma with which we started, it will be no comfort to the would-be anti-consequentialist if global consequentialism simply reiterates the simple act-consequentialist's claim that ascriptions of moral rights are false, but that (applying act-consequentialism to acts of speech) it is often morally right to speak such falsehoods.

To put the worry another way: the global consequentialist has told us what it is 'right' to say, but this notion of 'right' is potentially ambiguous. Clearly there is a distinction, in general, between some utterance's being *expedient* or *morally right to say* on the one hand, and its being *true* on the other. The question, then, is whether the notion of 'right to say' that is involved in global consequentialism is merely a matter of expediency, or one of truth; only in the latter case will global consequentialism have proceeded beyond simple act-consequentialism on the question of whether there are moral rights.

To have a distinctive view to consider, therefore, let us consider a version of global consequentialism that does indeed take on these more novel commitments. We then have a theory according to which the formula (GC) exhausts or otherwise determines all that can be said on moral matters. It must hold, then, either that there *is* no coherent question of which rights-ascriptions are true as opposed to merely expedient, or (slightly more concessively) that the global-consequentialist formula (GC) *determines truth-conditions as well as expediency-conditions*. According to this (latter) "thick" global consequentialism, the sentence "there is a moral right to life" is *true* if and only if the best moral system to (say) regulate one's behaviour with includes affirmation of a moral right to life. Insofar as there is indeterminacy in the choice of evaluative focal point and role – insofar, that is, as context fails to determine the relevant evaluative focal point ("moral system", or something else?)

and role ('regulating behaviour' in precisely what sense?), the question of whether or not there is a right to life *simpliciter* – whether or not the sentence “there is a right to life” is true – is itself correspondingly indeterminate (because relative to underspecified parameters). According to this “thick” version of global consequentialism, then, while questions of truth and expediency of course in general come apart (whether your haircut is nice is quite a different question, for instance, from that of whether my saying your haircut is nice would lead to good consequences), in certain normative domains they coincide.

This last claim is fairly radical, and will not be popular with primitive rights-theorists. The latter will want to insist that it would, or could, remain fundamentally *true* that there is a moral right to life even if empirical conditions were such that it was in no way expedient to integrate ascription of that right into any system, for any purpose. But disagreement with *primitive* rights-theory, of course, is to be expected: a significant part of the point of the whole exercise, recall, was to move away from that perspective on rights. By the lights of those who were convinced by the case for seeking a middle ground between primitive rights-theory and simple act-consequentialism in the first place, global consequentialism seems promising.

Two other issues. Before proceeding to apply global consequentialism to the case of war, I pause to flag two tangential issues.

First: any consequentialist theory presupposes some ‘theory of the good’, that is, a ranking of states of affairs (and hence of sets of consequences) in terms of better and worse overall. It is only relative to such a theory of the good that we know what we mean by one (say) act’s having “better consequences” than another. There are of course numerous controversial issues involved in the question of what is the correct theory of the good: in particular, it is controversial both which aspects of a state of affairs make it better or worse *ceteris paribus* (for example: just the well-being of sentient creatures, or other things too?), and how those aspects that each make states of affairs better *when other things are equal* are to be traded off against one another (for example, to settle comparisons between a state of affairs that involves more human well-being but also more damaged natural ecosystems, or comparisons that involve trade-offs of the well-being of some persons against that of others). There is a rich literature investigating issues of both of these types in detail, and it is not my purpose here to enter into the details. The present discussion is instead about how, in structural terms, the notion of overall good fits into moral theory: a matter about which act-, rule-, global- and non-consequentialist theories disagree, even when they agree on the theory of the good itself. Readers are free to interpret the discussion in terms of their own preferred theory of the good.³

³ Global consequentialism does, however, require, a ‘non-moral’ theory of the good – at least in spirit, if not in letter. That is, the theory of the good against which (say) candidate moral rules are to be evaluated should not already take account of issues of (say) desert, since questions of what a given individual *deserves* themselves depend on matters of which moral rules (if any) the individual in question has violated: it should depend in a more straightforward manner on (say) the well-being levels of individuals. In all likelihood, at the end of the day the global consequentialist will agree with the common-sense verdict that there is *some* sense in which it is ‘worse’ if an innocent is mistakenly imprisoned than if a genuine criminal is imprisoned; but that sense of ‘worse’ is not the one involved in the theory of the good that lies at the foundations of a global consequentialist approach.

Secondly: there is a potential worry about internal conflicts (of a sort) within global consequentialism. Any given action (for instance) falls with equal accuracy under a number of different descriptions, and in general, global consequentialism might deem one and the same action right relative to one accurate description, but wrong relative to another. Suppose, for instance, one finds oneself in a situation such that the right *act* (by global-consequentialist lights) is not the one that would be picked by following the right *decision procedure* (by global-consequentialist lights), nor the one that is mandated by the right laws (by global-consequentialist lights). What then ought one to do – perform the right act, follow the right decision procedure, or conform to the right laws?

I flag this worry only to set it aside (although the general issue will rear its head again in Section 4, below). The problem of conflicts may or may not, in the end, be sufficient ground for rejecting global consequentialism, and it is beyond the scope of this chapter to try to settle that question (for further discussion, see (Adams, 1976; Parfit, 1984, chapter 1; Railton, 1988; Feldman, 1993; Ord, 2008)). For the purposes of this chapter, I will simply assume that global consequentialism is promising enough for its implications to be worth working out.

Let us see, then, what global consequentialism has to offer by way of commentary on one recent first-order normative discussion. Section 3 takes a step back from the foundational issues that have been the focus of the present section, and reviews a recent discussion of the ethics and laws of war. Section 4 will turn to the task of exploring the connections between this recent discussion on the one hand, and global consequentialism on the other.

3. McMahan and Shue on the morality and laws of war

The laws of armed conflict. The existing international laws of war recognise the following three key principles:

- *Combatant equality*: While *political leaders* on one side of a war but not the other may be held to account for waging a war that is deemed unjust (that fails to satisfy the agreed principles of *jus ad bellum*), no legal distinction is drawn between “just combatants” and “unjust combatants”; combatants on both sides of a war are subject to the same rules concerning the conduct of their military operations.
- *Non-combatant immunity*: Combatants may only intentionally attack enemy *combatants*; non-combatants (for example, most civilians living in enemy territory) are not legitimate targets.
- *The privileged status of prisoners*: Prisoners of war may not be attacked or killed; they may only be detained for the duration of the war, to prevent them from serving again as combatants on the enemy side.

The orthodox view of the morality of war. The question then arises of what the relationship is between these *legal* principles on the one hand, and the *morality* of war on the other. According to the now-orthodox view (associated in particular with Michael Walzer (2006)), the above laws of war straightforwardly mirror the content of the underlying morality of war.

Walzer’s view proceeds from putative basic principles of morality concerning the notion of someone’s being (morally) “liable to be killed”. The idea is that (i) by default, one has a (moral) right

not to be attacked or killed, but (ii) one loses this right (becomes 'liable to be attacked/killed') if one poses a sufficiently serious (perhaps lethal) threat to others. This is how Walzer seeks to defend the view that in a context of war it is often morally permissible to use lethal force, while in a peacetime context killing is normally one of the most serious of wrongs: killings in war, but not in ordinary life, are generally carried out in settings in which the victim was posing a threat.⁴

Given these principles, the further principles of combatant equality, non-combatant immunity and the privileged status of prisoners follow immediately: just combatants, no less than unjust combatants, pose a lethal threat to their adversaries, and neither non-combatants nor prisoners of war *pose* any such threat (although non-combatants may sometimes *contribute* to threats, for example by supporting the army in various ways).

Criticisms of the orthodox view. Recently, however, there has emerged something of a consensus among moral theorists that Walzer's putative underlying principles of morality are simply not true, *qua* moral principles. 'Posing a threat', these revisionists argue, is neither necessary nor sufficient for becoming liable to be killed.

It is not sufficient, because: (1) If the threat you pose is *justified*, then you are not liable as a result of posing that threat. (For example, if you are defending yourself against a violent mugger, and the only mode of resistance available to you is the use of potentially lethal force, you do not lose your right not to be killed by the mugger, because your threat is a justified one – it is only the mugger who has lost his right not to be attacked and potentially killed, by posing an *unjustified* threat against you in the first place.) (2) If you are not *responsible* for the threat you pose, then you aren't as a result liable. (For example, if you are stuck at the bottom of a well and an innocent third party is thrown down the well, and will crush you to death on landing unless you detonate a bomb that will blow him up *en route*, he poses a lethal threat but (intuitively) is not *liable* to be killed as a result, although your killing him may be *excusable*.) In the context of war, the first point suggests (to a first approximation) that only combatants on the unjust side of the war, not those on the just side, are liable; the second point suggests that not all unjust combatants are liable, and perhaps not even most.

Actually *posing* the threat is also intuitively not *necessary* for becoming liable to be killed: if, for instance, you have hired a hitman to kill some third party, then intuitively, you (no less than the hitman himself) have thereby lost your right not to be attacked – you have become liable thanks to your *contribution to* and *responsibility for* the threat, even though you yourself are (arguably) not *posing* the threat. (In the context of war, the point would apply e.g. to the politicians who authorised an unjust war, and arguably to some ordinary civilians in virtue of their support of the war and of the armed forces.)

⁴Importantly, the condition that someone is 'liable to be killed' is neither necessary nor sufficient for its being *permissible* to kill him – liability to be killed implies only that the person in question has lost his *right* not to be killed, so that he is not *wronged* if he is killed. It would remain impermissible to kill him if, for example, no sufficiently good purpose would be served by such killing. But the idea is that often in a context of war, some sufficiently good purpose would be so served, so that considerations of liability on the one hand and those of the permissibility of acts of killing on the other very often line up with one another in practice.

Revisionist liability theory and the laws of war. Suppose that one tried to take an approach that was structurally like Walzer's – that is, to advocate laws of war that simply follow from a straightforward application of the underlying morality of war – but starting, not from Walzer's own (crude) moral theory, but from the revisionists' more sophisticated liability theory.

It is clear, in the first instance, that the resulting laws could not look anything like the existing laws of armed combat. The principle of combatant equality would have to go, since revisionist liability theory holds that only those who pose or contribute to an *unjust* threat are liable, and this will exclude many or most combatants on the *just* side, at least during many of their missions. And the principle of non-combatant immunity would also have to go, since revisionist liability theory holds that suitable *contributions* to an unjust threat are (if one is responsible for said contributions) sufficient for liability, and this condition will presumably be met by some non-combatants: for example, by the politicians on the unjust side who authorised an unjust war, and perhaps also e.g. to those civilians who manufacture arms for supply to the army. (For a detailed discussion of precisely which non-combatants would be liable on this view, see Fabre (2010).)

Secondly, however, given the radical uncertainties that soldiers and others will face concerning the questions of who is and is not liable – which side in the war is the just side, and which particular enemy combatants and noncombatants meet the conditions for liability – it seems that on any plausible approach to moral permissibility under conditions of uncertainty, the results of applying the revisionist liability theory will be that just about *any* real-world lethal attack will be morally impermissible ('contingent pacifism'). If so, then the corresponding laws would have to deem all actual wars illegal.

Theorists of the ethics of war thus face the following trilemma: First, they could simply bite the bullet, and advocate pacifism at the level of the laws of war, no less than at the level of its morality: that is, they could advocate simply outlawing all war. Second, they could hold that there is a sharp distinction between the morality of war and its laws, even its morally optimal laws: on this view, while the *morality* of war is as given by the revisionists' liability theory, there are good reasons why the morally optimal *laws* will look a lot more like the existing laws of armed conflict. Or, third, they could reject the revisionists' liability theory as an account of the morality of war.

The first horn of this trilemma seems implausible, and I am not aware of anyone who has advocated it. The second horn is grasped by Jeff McMahan; the third by Henry Shue. I turn now to a brief exposition of McMahan's and Shue's views, in turn.

McMahan's view. According to McMahan, then, the *morality* of war does indeed track the revisionists' liability theory. However, we are to distinguish sharply between the morality of war on the one hand, and the laws of war on the other; according to McMahan, there is no reason why these two codes should be particularly close in content, and indeed the existing laws of war may (for all McMahan says) be very close to the morally optimal ones.

Since laws are made by conventional agreement, McMahan agrees that the question of which are the morally optimal *laws* is highly sensitive to considerations of what the practical consequences would be of enacting this or that law. And there are many reasons for thinking that any attempt to enact the revisionists' liability principles directly into law – in the way envisaged as the "first horn" of

the trilemma above – would have disastrous consequences. In the first instance, epistemic limitations are important: since the laws of war must be action-guiding, they cannot draw distinctions that soldiers in the heat of battle would not be able to draw. Thus, while they can distinguish between enemy combatants and enemy non-combatants as broad categories, and between friendly and enemy combatants as broad categories, they cannot distinguish between liable vs non-liable individual combatants or non-combatants, nor between the just vs the unjust side of the war. Secondly, there are also issues of bad faith: laws according to which any killing of “just” by “unjust” combatants is illegal would be open to abuse by victorious parties seeking revenge, who will tend to simply declare their own side of the war to have been the just one, regardless of where the moral truth lies. Thirdly, if (as I tentatively suggested above) the consequence of insisting that the laws must closely reflect revisionist liability theory would be a legal contingent pacifism, according to which all actual wars are illegal, the laws would in practice simply be disregarded; better, surely, to have laws that acknowledge the reality of war, and do something to try to limit the corresponding damage.

According to McMahan, however, since the *morality* of war (unlike its laws) is *not* a matter of convention, it just “is what it is” (2008, p.35), and pragmatic concerns are here irrelevant. It is, on this view, simply a moral *fact* that just combatants typically have a right not to be killed, notwithstanding the inadvisability of trying to take any account of this fact when devising laws.

The central worry for this view is how to handle the conflicts that (on this view) can arise between morality on the one hand, and law on the other. Each source of normativity (moral and legal) might say, of a given action, that that action is either required, permitted (but not required) or forbidden; there is no general guarantee that the two sources will issue the same verdict in any given case.

McMahan is relatively unconcerned by cases in which morality permits some action that the law either forbids or requires, or vice versa: in these cases, he says, one ought to obey the prohibition or requirement, setting aside the mere permission.⁵ But conflicts between a requirement from one source and a prohibition from the other are more serious. What *ought one to do*, that is, if and when morality requires some action that the law forbids, or vice versa?

While McMahan thinks that such conflicts are likely to be relatively rare in practice, he concedes that they can arise in principle. (For example, he thinks that one might sometimes be morally *required* (not just permitted) intentionally to target liable civilians, when this is required for the success of one’s just mission (ibid., p.38).) His verdict is simply that in any and all such cases, one ought to follow the dictates of morality over those of law (while openly acknowledging the legal violation one thereby commits, but pleading leniency on the grounds of one’s “moral justification” (ibid, p.39) for breaking the law).

Shue’s view. Shue, meanwhile, claims to find McMahan’s position unintelligible. According to Shue, while there is always an open question of whether or not one ought to obey *actual* laws – those laws might themselves be immoral – it is always the case that “one morally ought to obey the morally best laws” (2008, p.110), in any context in which needs laws at all. If so, then there is no question of having a separate morality that can sometimes conflict with (the morally best) laws, as proposed by

⁵ For dissent even on this count, see Shue (2008, p.107).

McMahan; the task is simply to work out what the morally best laws are, and morality will then be exhausted by the injunction to obey those laws.

The key distinction for Shue is thus not between the morality of war on one hand and the laws of war on the other; rather, it is between what morality requires *in the context of ordinary life* on the one hand, and what morality requires *in a context of war* on the other. Shue advocates a picture according to which all of morality – for ordinary life, for war, or for any other context – is grounded in an unchanging set of “fundamental moral considerations” (ibid, pp.88, 90, 95), for instance, principles that are in favour of well-being, and against cruelty. The matter of which “specific standards” follow from those fundamental considerations, however, can depend on the context: in particular, it could well be that *in a context of ordinary life* a fundamental concern for welfare leads to a specific standard that simply bans (just about) all killing, while *in the context of war* those same fundamental concerns instead lead to specific standards that merely issue guidance on whom one is vs is not permitted to kill. According to Shue, McMahan’s makes a basic mistake in applying the *specific standards* appropriate to ordinary life to the very different context of war.

4. Global consequentialism meets the morality of war

What should a global consequentialist make of the exchange between McMahan and Shue? I will argue here that by global consequentialist lights, there are some particular and important respects in which McMahan is correct and Shue incorrect, others in which their status is reversed. And it is, in purely general terms, highly plausible that some first-order normative disagreements might be traceable to disagreements at some more foundational level. The possibility of global consequentialism as at least *one*, arguably somewhat appealing, account of the foundations of morality therefore has some illumination to offer on the apparent stalemate between these two theorists. Six points are worthy of note.

First: By global consequentialist lights, McMahan overstates the difference between the morality vs. the laws of war. Recall that McMahan thinks that it follows, simply from the fact that morality is not the product of any *convention* or *agreement*, that considerations of ‘reasons for adopting’ one or another moral standard, such as the pragmatic considerations that he allows centre stage in the discussion of the *laws* of war, cannot have any role to play in determining the content of the *morality* of war.⁶ But if global consequentialism is correct, this is not quite right: instead, in the domain of morality no less than in that of law, considerations of (at least) the *consequences of one*

⁶ Thus, for example, McMahan writes: “There is today a broad tendency to discuss the ‘rules of war’ without indicating whether these rules are to be understood as moral rules, legal rules, or both. Yet, it is often also suggested that these rules are not discovered but are instead designed to serve certain purposes. George Fletcher, for example, writes in defence of the equality of combatants that ‘the *reason for adopting* a rigorous distinction between jus ad bellum and jus in bello is the need for a bright-line cleavage that is workable in the field of battle. ...’ And Larry May contends that ‘the rules of war are *designed* to help societies, and States, meet their responsibilities towards those [their own combatants] who have had their vulnerability increased by that same society, or State.’ Many of those who hold that the rules of war are thus artefacts devised by human beings to serve certain functions do not take the same view of the principles of morality. Nor should they. They thus implicitly acknowledge a distinction between basic, non-conventional moral principles and conventional or legal principles, both of which may simultaneously govern the practice of war” (McMahan 2008, pp.34-5; italics in original.). In this passage, McMahan apparently takes it that Fletcher’s and May’s use of locutions like “reason for adopting” and “designed” commits them to the view that the rules under discussion are matters of agreement or convention, and hence are legal rather than moral rules.

or another moral standard being advocated/adopted/used as a decision procedure/used as the basis for ascriptions of responsibility and blame/etc. are highly relevant to the question of which moral statements are true (i.e., to the content of morality). The admitted fact that morality is not a matter of *convention* does not, on this view, prevent considerations of consequences from playing a role in determining its content.

Second: As we have seen, one of Shue's key objections to McMahan is that Shue regards the very idea of a morality of war that diverges from the prescriptions of the (morally best) laws of war as incoherent. But if global consequentialism is correct, then there is a clear route by which this divergence could easily happen. This is because according to global consequentialism, the question of which are the right standards is intelligible only relative to a particular *role* those standards might play, and morality corresponds to a different role (really, a set of roles) than does law. It could easily be the case, for instance, that the principles whose *enaction as laws* has the best consequences (hence, the content of the 'morally best laws') are quite different from the principles whose *use as a decision procedure by conscientious moral agents* has the best consequences, since considerations of transparency and enforceability are relevant to the former but not the latter. On this point, by global consequentialist lights, Shue is incorrect and McMahan in the end correct, although not for the reason McMahan himself gives.

Third: Granted that the content of the morality of war could differ even from that of the morally best laws for war, however, the question remains of *how* different its content will be. McMahan, we have seen, thinks that the content of the relevant morality is *very* different from that of the relevant laws. For most of the 'roles' relative to which we might (as global consequentialists) assess a system of morality, however, such radical difference is unlikely to be the result of applying the global-consequentialist formula (GC). The basic reason for this lies in observations already made by McMahan himself: for example, that "[c]ombatants should be reluctant to give their individual judgment [about the permissibility of killing based on judgments of individual liability] priority over the law, for the law has been designed in part precisely to obviate the need for resort to individual moral judgment in conditions that are highly un conducive to rational reflection" (ibid, p.41). For most *moral* purposes ('roles'), be they ones of individuals' decision procedure or ascriptions of moral responsibility, the result of such observations (according to global consequentialism) is not that combatants should defer to the law rather than to morality, but rather that the content of morality of war itself is much closer to that of law than it would be if the context of war involved more ideal epistemic conditions. Thus, while McMahan is correct in his qualitative claim that the content of the morality of war might well diverge from that of the (morally best) laws of war, he significantly overstates the extent of the divergence.

Fourth: While McMahan thinks that morality is fairly unitary (that the morality of war is just the morality of ordinary life, applied to the context of war, but with the same principles of e.g. conditions for liability applicable in both contexts), Shue sketches the outline of a more complex picture, according to which context-invariant "fundamental moral concerns" give rise to differing "specific standards" in different contexts. However, Shue says nothing further about *how* the fundamental concerns he gestures at might "give rise" to diverging specific standards. Global consequentialism offers one principled way of filling in the picture, and one that (moreover) seems much within the spirit of what Shue does say. Shue's "fundamental moral concerns" become a

comprehensive theory of the good (so that we know what counts as the best overall consequences), and the notion of this theory of the good “giving rise” to a particular specific standard is explicated, relative to a role for the standard in question, by the global-consequentialist formula (GC). On this point, then, by global consequentialist lights McMahan is incorrect and Shue correct, but global consequentialism is also able to fill in the crucial details where Shue’s own discussion leaves matters extremely vague.

Fifth, a minor observation on the structure of the landscape. Commenting on McMahan, Rodin writes that McMahan takes the morality of war to be “predominantly deontological in character”, while the “morality of public rule formation” is “predominantly consequentialist in character” (2011, p.453). This characterisation of McMahan’s views, however, is potentially seriously misleading. For McMahan does not think that either the morality of war *or* the “morality of public rule formation” (the morally best laws for war) is consequentialist *in content*. (Indeed, no-one has ever advocated, nor would they ever advocate, laws that are genuinely consequentialist *in content*: a law according to which one is simply legally required to do whatever would promote the best overall consequences would be a laughing stock, and rightly so.) The point is rather that, as noted above, McMahan (like the global consequentialist) takes consequentialist considerations to be *relevant to determining* the content of the laws of war, but (in contrast to the global consequentialist) takes such considerations to be irrelevant to determining the content of the morality of war. This distinction is crucial: failure to recognise it seems likely to block all possibility of finding any middle ground between the extremes of an overly simplistic act-consequentialism and an overly primitive rights-theory.

Sixth, and finally: We saw that, as a result of admitting that the content of the morality of law might diverge from that of even the morally best laws for war, McMahan was compelled to acknowledge the possibility of conflicts between requirements and prohibitions issuing from the two different sources. This problem is arguably simply the reappearance, in the particular context of war, of the problem of conflicts that plagues global consequentialism quite generally, and that I outlined in Section 3. To note that the problem is a general one is of course not immediately to say anything about how it is best to be solved, but it may nevertheless prove illuminating.⁷

I note in closing that it is also true that nothing about the context of war exorcises the purely general worries that we noted in Section 2 for global consequentialism. In particular, those who are dissatisfied with global consequentialism’s partial collapsing of truth-conditions and expediency-conditions in general will be no more happy with this approach in the application to war (for some complaints that are easily interpreted along these lines, see e.g. Rodin (2011, pp.459-61). Global consequentialism is far from uncontroversial. But it is also far from implausible.

Acknowledgements

⁷ Further, the spirit of global consequentialism perhaps also suggests a principled strategy for dealing with such conflicts, via ascent to a meta-level: “the right approach to take towards conflict-resolution is...”.

For valuable discussions, I am grateful to Henry Shue, to the members of the *Human Rights for Future Generations* Project and to the participants in a workshop held in Oxford in May 2016, especially Dapo Akande and David Rodin.

Bibliography

- Adams, R. M. (1976). Motive utilitarianism. *The Journal of Philosophy*, 73(14), 467-481.
- Crisp, R. (1992). 'Utilitarianism and the life of virtue', *Philosophical Quarterly* 42:139–160.
- Fabre, C. (2010). 'Guns, food, and liability to attack in war', *Ethics* 120: 36-63.
- Feldman, F. (1986). *Doing the best we can*, (Dordrecht: Reidel).
- Feldman, F. (1993). 'On the consistency of act- and motive-utilitarianism: a reply to Robert Adams'. *Philosophical Studies* 70:201–212.
- Hooker, B. (2000). *Ideal code, real world: A rule-consequentialist theory of morality*. Oxford University Press.
- Kagan (1992), 'The structure of normative ethics', *Ethics* 6:223-42.
- Kagan, S. 2000. 'Evaluative focal points', in Hooker, Mason and Miller (2000), pp. 134–155.
- Mason, E., Hooker, B., & Miller, D. E. (2000). *Morality, rules, and consequences: a critical reader*. Rowman & Littlefield.
- McMahan, J. (2008). 'The morality of war and the law of war', in Rodin and Shue (2008), pp.19-43.
- Narveson, J. (1973). Moral problems of population. *The Monist*, 62-86.
- Ord, T. (2008). Beyond action: Applying consequentialism to decision making and motivation. DPhil thesis, University of Oxford.
- Parfit, D. (1984). *Reasons and persons*. OUP Oxford.
- Pettit, P. and Smith, M. (2000). 'Global consequentialism.' In Hooker, Mason and Miller (2000), pp. 121–133.
- Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *Philosophy & Public Affairs*, 134-171.
- Railton, P. (1988). 'How thinking about character and utilitarianism might lead to rethinking the character of utilitarianism', *Midwest Studies In Philosophy* 13:398–416.
- Rodin, D. and H. Shue (2008). *Just and unjust warriors: The moral and legal status of soldiers* (eds.). Oxford University Press.
- Rodin, D. (2011). 'Morality and law in war', in H. Strachan and S. Scheipers (eds.), *The changing character of war*, pp. 446-63, Oxford University Press 2011.
- Shue, H. (2008). 'Do we need a 'morality of war'?' In Rodin and Shue (2008), pp.87-111.
- Walzer, M. (2006). *Just and unjust wars: A moral argument with historical illustrations*, 4th edition. New York: Basic Books.

Williams (1973). 'A critique of utilitarianism', in Smart and Williams, 'Utilitarianism: For and against'.