

NON-COMBATANT IMMUNITY AND ‘JUST WAR’.

Amy Russell

Victoria University of Wellington

Moral philosophers and the international political community alike have traditionally valued the lives of civilians above those of soldiers. This thesis asks why this should be so, and suggests that the only answer should lead us to re-examine key beliefs about the morality of war.

Acknowledgements

I am grateful to Ramon Das for his outstanding supervision; to Russell Brown for his stimulating input; to my parents for their continual support, material and otherwise; and to all the graduate students of the VUW Philosophy Department (especially Ben Brooks, Vanessa Schouten, Dennis Poole, and Alan Poole) for their imaginative suggestions, which did much to enliven my research time. I am also grateful to the administrators of the Herbert Sutcliffe scholarship for their very welcome financial support.

Above all, I am grateful to Allie Jarratt, whose unstinting optimism and generous friendship immeasurably improved my enjoyment of this project.

CONTENTS

Introduction	4
Overview of chapters	5
Chapter 1: Laying the groundwork	8
Section 1: Establishing terms of reference	8
Section 2: Our intuitions about killing	9
Section 3: The interaction between <i>jus ad bellum</i> and <i>jus in bello</i> and the mutual entitlement to kill	10
Chapter 2: Innocence and guilt	21
Section 1: A definition of innocence	21
Section 2: Deciding the guilt of soldiers and the innocence of civilians	25
Section 3: Justifying the killing of innocent soldiers	32
Chapter 3: Justifying self defence	39
Section 1: Necessary and sufficient conditions for self defence	40
Section 2: Explaining deontological defensive permissions	44
Section 3: The utility of ‘either him or me’ killing	48
Section 4: Self defence in law	54
Section 5: Conclusions	57
Chapter 4: Self defence and war	59
Section 1: The conditions for self defence	60
Section 2: Defending the principle of non-combatant immunity via self defence	64
Section 3: The paradigm combatant as a non-innocent threat	70
Section 4: Conclusions as regards immunity theory	76
Chapter 5: Policy and moral conventions	78
Section 1: Convention-dependent morality and utilitarian heuristics	78
Section 2: Putting non-combatant immunity to the test	83
Section 3: Ideal vs. non-ideal theory	94
Chapter 6: Concluding remarks	104
Bibliography	111

INTRODUCTION

Moral philosophers and the international political community alike have traditionally valued the lives of civilians over those of soldiers. The first part of *jus in bello*, the doctrine which aims to characterise the just conduct of war, states that ‘civilians, as non-combatants, must not be attacked or killed’, whereas the only requirement concerning the killing of soldiers is that any attack must meet the requirement of proportionality: it must not cause so much harm that the good it does is overridden.¹ Similarly, Article 51 of the Geneva Protocols states that ‘the civilian population and individual civilians shall enjoy general protection against dangers arising from military operations’, and that ‘the civilian population as such, as well as individual civilians, shall not be the object of attack’.² The requirement of proportionality is mentioned only with reference to the protection of civilian life or cultural objects, except in the general statement that ‘it is prohibited to employ weapons, projectiles and material and methods of warfare of a nature to cause superfluous injury or unnecessary suffering.’³ The specific protections offered to combatants are limited to wounded, sick or shipwrecked combatants, and prisoners of war – those combatants who most closely resemble civilians. The Protocols do state that all attacks must be limited to ‘military objectives’, but the definition of these objectives is permissive, to say the least:

Military objectives are limited to those objects which by their nature, location, purpose or use make an effective contribution to military action and whose total or partial destruction, capture or neutralization, in the circumstances ruling at the time, offers a definite military advantage.⁴

¹ Quoted from Norman, Richard, *Ethics, Killing and War* (Cambridge: Cambridge University Press, 1995), p.119

² Sections 1-3 of Article 51 of the Geneva Protocol (additional to the Conventions of 1949) relating to the Protection of Victims of International Armed Conflicts (Protocol I, 8 June 1977).

³ Section 2 of Article 35 of the Geneva Protocol (additional to the Conventions of 1949) relating to the Protection of Victims of International Armed Conflicts (Protocol I, 8 June 1977). What counts as ‘unnecessary’, and therefore illegal, suffering in the context of war is a difficult question. The standard of military necessity, which might be thought a promising candidate, is mentioned in the Conventions only as a justification for harm caused to civilians.

⁴ Section 2 of Article 52 of the Geneva Protocol (additional to the Conventions of 1949) relating to the Protection of Victims of International Armed Conflicts (Protocol I, 8 June 1977).

To kill enemy soldiers in large numbers surely offers a definite military advantage.

This thesis examines the moral basis for the distinction that these laws and doctrines draw between soldiers and civilians. I explain why the distinction between combatant and non-combatant casualties is not, in a significant proportion of cases, a morally sound one. I argue that any moral justification of the principle of non-combatant immunity must be of a utilitarian nature, pointing to its ability to limit the overall carnage of warfare. The implications for *jus in bello* of recognising that the principle can be justified only on these grounds are wide-ranging and important. If we want to retain civilian immunity, we must accept a utilitarian simulacrum of that doctrine. I argue that applying utilitarian standards to the just conduct of war will lead us to prefer very different sorts of policies from those currently embodied by *jus in bello*. Thus what we think about civilian immunity may have consequences for what we think about the moral foundation of our doctrine of just war.

Overview of chapters

In chapter 1 I lay the groundwork for my later discussion, identifying the importance of our intuitions about the wrongness of killing, and about when it is sometimes justified. From these I introduce two concepts as promising justifications for the principle of non-combatant immunity: protecting the innocent, and killing in self defence. I also spend some time explaining why I think that the mutual entitlement to kill ought to be retained as part of *jus in bello* despite the counterintuitiveness of according equal permissions to conflicting warring parties.

In chapter 2 I discuss the first set of arguments in defence of the principle of non-combatant immunity: arguments from innocence. I first argue for an understanding of innocence as an important moral notion distinct from that of self defence. I then argue, using some clear problem cases, for the conclusion that even if an agent's innocence or guilt does give us a reason to protect or to target him, this moral consideration cannot support a distinction between combatants and non-combatants.

Chapters 3 and 4 are devoted to the difficult question of killing in self defence. I discuss in chapter 3 how self defence may be justified in the personal case, with a view to better understanding the nature and extent of our defensive permissions. Then in chapter 4 I apply my findings to war, arguing that the analogy between personal self defence and defensive killing in war will often be unsuccessful. I then argue that, just as with arguments from innocence, although self defence may give us a reason to target or to protect certain people, it cannot give us a reason to distinguish combatants from non-combatants. At the end of chapter 4 I summarise the uncomfortable position of the immunity theorist, given that arguments from innocence and from self defence do not succeed in justifying the principle of non-combatant immunity. I suggest that any justification for the principle must be found extrinsically in the role it plays, rather than in the intrinsic qualities of the persons between whom it discriminates.

In chapter 5 I consider this possibility of extrinsic justification for non-combatant immunity. I discuss the idea of convention-dependent morality and argue that conventions and policies in war ought to be adopted or rejected according to the utilitarian consideration of how well they prevent harm. I argue further that when compared to the feasible (in the near future) alternatives, non-combatant immunity can be defended as a convention which is highly successful in controlling the harms of war, and for this reason ought to be retained as part of *jus in bello*, at least in the immediate future. I move on to discuss the dual task of the political philosopher, in particular finding a balance between ideal and non-ideal theory. I argue that although we have an obligation to adhere to non-combatant immunity in the short term, there is no reason to think that this policy will guarantee moral behaviour in the long term.

In my final chapter, I relate my conclusions from chapter 5 to my broad conclusions about non-combatant immunity as a moral and political doctrine. I argue that there are three central questions which remain to be answered before the extent of our obligations concerning just war become clear: ‘What sorts of policies or conventions might be implemented in the future?’; ‘How we are to weigh their merits?’; and ‘What political and social changes are needed before they stand any chance of success?’ I argue that a proper recognition of the utilitarian nature of non-combatant immunity will greatly change the importance of the notion of military necessity in the just conduct of war, making this notion the primary standard of acceptable killing:

Just as we currently think that civilian deaths, if they are to be justified, should be demonstrably necessary to the successful conduct of war, so we should demand that soldiers' deaths be treated as a serious part of the harm of war, to be avoided where possible. I conclude that a proper understanding of the harms of war requires an understanding of the harms of combatant deaths, and that the widespread recognition of this fact will be an important step toward establishing a political world committed to punishing violence and protecting peace.

CHAPTER ONE

Laying the groundwork

Section 1: Establishing terms of reference

In this chapter I hope to lay the groundwork for a thorough discussion of non-combatant immunity. In order to do this, I shall make several stipulations of a general nature.

First, from time to time I will make reference to common law to illustrate the moral implications of some of the arguments I discuss. In the majority of cases my references will be internationally relevant; where not, it is New Zealand law to which I refer.

Second, throughout my discussion I use the terms ‘soldier’ and ‘combatant’ interchangeably, and likewise ‘civilian’ and ‘non-combatant’. I am aware that there can be a distinction drawn between ‘soldiers’ who fight and ‘combatants’ who are part of the army but do not fight, but it is the difficulties inherent in this distinction which I particularly want to discuss. For the purposes of argument about ‘non-combatant immunity’, therefore, I will assume it is the same thing as ‘civilian immunity’. In chapters 2, 3 and 4 I discuss at length how these terms may be refined, arguing that any standard of identity appealed to will either exclude those the immunity theorist wishes to include, or include those she wishes to exclude.

Third, in the arguments below, and particularly in chapters 2 and 3, I make much reference to our widely held intuitions about killing. It is important to note that I do not think appeal to intuitions is a substitute for moral argument. But I do think that identifying our intuitions plays an important part in explaining our general moral beliefs. I discuss this further in section 2, below, and also in section 3 of chapter 3.

Fourth, my normative conclusions throughout this discussion are designed to be applied to both sides of a conflict, regardless of the justice of their *ad bellum* cause. In section 3 of this chapter I give an argument for this, and for accepting a mutual entitlement to kill as part of *jus in bello*.

Section 2: Our intuitions about killing

Our general intuition about killing can be phrased as something like ‘killing is usually wrong because it is usually wrong to deprive someone of their life’. In other words, killing is usually wrong simply because it *is* killing. This intuition can be explicated with reference to a number of different moral theories: rights theorists will talk of the right to life, virtue ethicists might invoke the virtue of acting peacefully, and utilitarians can point to the negative consequences of acts of killing. I will not go into a long discussion here about which of the competing moral theories most successfully captures our intuitions about killing. At this point it is enough to say that the intuition that it is usually wrong to kill innocent people is universally shared throughout the moral community; and any theory which cannot accommodate this much fails as a moral theory.

In the following chapters I often refer to our intuitions about when it is and is not right to kill in the case of soldiers and civilians. As I have stated above, and as my later discussion will demonstrate, this is not because I think these intuitions do not require, or do not allow of, explicit justification by some moral theory. But I believe it is instructive to plainly identify our intuitions about killing because, as I argue, it is on the intuitive level that the combatant/non-combatant distinction has its greatest appeal.

The limits on killing

As well as a general agreement that killing is usually wrong, we share reasonably general intuitions about at least some of the sorts of cases in which killing may be justified, for example when someone is killed in self defence, or when the person killed is ‘non-innocent’ in some important moral sense (as in capital punishment).

The history of civilian immunity has traditionally appealed to one of these two notions, self defence or innocence, to justify the distinction between civilians and soldiers. The terms ‘innocence’, ‘guilt’ and ‘self defence’ are imprecise, however, and I spend chapters 2, 3 and 4 elucidating them and evaluating the arguments which use them to defend a principle of non-combatant immunity.

Section 3: The interaction between *jus ad bellum* and *jus in bello* and the mutual entitlement to kill.

In what follows, I make a number of assumptions about the possibility of just warfare. This is not because I am convinced that actual wars can be or sometimes are fought with ‘just cause’ as outlined in *jus ad bellum*. It is rather because I believe moral philosophy ought to be action-guiding, and that it must therefore make allowances for the reality of international politics. I consider this issue more closely in chapter 5. Here I want to briefly discuss one result of the influence of real-world necessity over moral theory: justification for the strangely counterintuitive interaction between *jus ad bellum* and *jus in bello* and for the latter doctrine’s endorsement of the mutual entitlement to kill. As elsewhere in this essay, I argue that this feature of Just War theory can only be justified on utilitarian grounds. In other words, it is right that soldiers have a mutual entitlement to kill only because the absence of such an entitlement is consequentially undesirable.

The puzzle

In *War and Self-Defense*, David Rodin identifies what he considers a strange feature of Just War theory:

If an aggressive war is fought within the bounds of *jus in bello*, then the Just War Theory is committed to the seemingly paradoxical position that the war taken as a whole is a crime, yet that each of the individual acts which together constitute the aggressive war are entirely lawful.⁵

⁵ Rodin, p. 167

This ‘seemingly paradoxical position’ is a result of the separation of *jus ad bellum* from *jus in bello*, in which the moral edicts of the latter make no reference to those of the former. Thus the moral status of a warring state as regards *jus ad bellum* has no effect on that state’s warring rights under *jus in bello*. Rodin again:

Jus ad bellum begins from the assumption that war is an illegal activity whose legitimate use is limited to, at most, one side of a given dispute. *Jus in bello*, on the other hand, assumes that war is a law-governed and hence implicitly lawful activity with entitlements and restrictions that accrue *equally to both sides*.⁶

So *jus ad bellum* tells states when they may go to war, and *jus in bello* tells them how they may act in war. The puzzle that Rodin identifies and discusses is this: we know that, in any given dispute, *at most one* side (and more usually neither side) has complied with *jus ad bellum*. Thus at most one side is justified in being at war. How can it be, then, that both sides are judged to be acting morally rightly in war, provided they adhere to the rules of *jus in bello*? Why do we judge each part of the war endeavour separately, and more specifically, why should we believe that soldiers on each side of a war possess a mutual entitlement to kill, as specified in *jus in bello*?⁷ Whence could this entitlement derive?

It is very important that these questions have a good answer, because the content of *jus in bello* is hugely affected by the requirement that it apply equally to both sides. If we were to devise a set of rules of warfare to apply only to those fighting just wars, they may well be greatly more permissive than *jus in bello* – and we would need to devise a second, greatly *less* permissive, set of principles for the unjust side to follow, which would likely be restrictive to an extreme degree. So why not endorse this dual conception of *jus in bello*?

⁶ Rodin, p. 166, italics added.

⁷ The first principle of *jus in bello*, that of discrimination, identifies those who may not be attacked, including civilians, prisoners of war, and soldiers who surrender. Attacks on the enemy’s active soldiery are unlimited (except by concerns of proportionality as stipulated in the doctrine’s second principle). Thus every soldier is permitted to kill enemy soldiers, regardless of the justice of his cause.

As I indicated above, my answer to this question is intimately linked to the regrettable facts of real-world war. I argue below that *jus in bello*'s moral symmetry between warring sides, including the 'mutual entitlement to kill', can be justified only because removing this symmetry would be consequentially undesirable – and were it not so, no such approach to *jus in bello* would be defensible. Thus the limited moral interaction written into *jus ad bellum* and *jus in bello* can only be justified by an appeal to utilitarian considerations (and I argue in the following chapters that this is also true for *jus in bello*'s requirement of non-combatant immunity). But some writers, notably Michael Walzer, have tried to ground *jus in bello*'s mutual entitlement to kill in more foundational or deontologically relevant moral notions such as consent and duress.⁸ I outline his position below. My discussion of Walzer closely follows Rodin's in chapter 8 of *War and Self-Defense*, though the conclusions we draw differ greatly.

Whence the mutual entitlement to kill?

A commonly offered⁹ reason why soldiers may be targeted in war is that they have consented to soldiery, and part and parcel of soldiery is a liability to be killed by the enemy. If this were the case, it would seem reasonable to allow a mutual entitlement to kill, for consider other situations in which people seem to possess a permission to harm one another: duellists fighting for honour, perhaps, or Rodin's example, boxers in the ring. In these cases, the fact that the persons harmed consented to their situation does justify the harm inflicted upon them, provided that the rules mutually agreed to governing the harmful behaviour are adhered to. But can the same justification apply to soldiers killing one another in war? It would seem not. First, the claim that all soldiers consent to soldiery is certainly a false one. As I discuss in chapter 2, the required conditions for consent are often absent.¹⁰ But even though some soldiers do consent to engage in warfare they know to be dangerous, this is far from consenting to be killed by the enemy. Merely acknowledging a risk is not consenting to being exposed to it, just as my acknowledgment that I may be hit by a truck tomorrow does not amount to my consent to that situation. In Rodin's words:

⁸ Walzer, M., *Just and Unjust wars* 2nd ed (New York: Basic Books, 1992), pp. 34-47 and *passim*.

⁹ By entry-level Philosophy students

¹⁰ See section 2 of chapter 2.

Soldiers fight knowing they may be killed, but they do not thereby permit their enemies to kill them; they do not say to their enemies 'you may kill me if you can'.¹¹

So the claim that the mutual consent of soldiers constitutes a mutual entitlement to kill is false. Michael Walzer, however, uses this notion of consent differently, and offers an opposite explanation for why soldiers have an equal entitlement to kill. Rather than arguing that soldiers mutually consent to be killed, Walzer argues that the equal entitlement to kill derives from the fact that *neither* side's soldiers have consented to be killed. He argues that to give soldiers an equal right to kill the enemy is one of the few ways to limit the terrible injustices of war:

War is still, somehow, a rule-governed activity, a world of permission and prohibitions. . . there is a license for soldiers, and they hold it without regard to which side they are on; it is the first and most important of their war rights. They are entitled to kill, *not anyone*, but men whom we know to be victims. We could hardly understand such a title if we did not recognize that they are victims too.¹²

Without the equal right to kill, war as a rule-governed activity would disappear and be replaced by crime and punishment, by evil conspiracies and military law enforcement.¹³

Walzer's claim, at least in the first of these quotations, seems to be that because all soldiers are equally victims of the base machinations of their state, they are equally allowed to kill one another. However, it would seem to make greater sense to urge the opposite claim: that soldiers are equally *prohibited* from killing one another, as they can recognise in their enemy the illegitimate target they themselves represent. The mutual lack of consent of soldiers would seem to forbid, rather than to allow, mutual killing, and Walzer's opposite claim stands in need of strong argument.

¹¹ Rodin, p. 172

¹² Walzer, p. 36. See also discussion in Rodin, pp. 170-172.

¹³ Walzer, p. 41. Why 'military law-enforcement' is necessarily worse than innocent victims possessing a right to kill one another is not clear to me; surely it must depend on the justice of the laws and the integrity of the enforcement.

Likewise in need of support is the implication, in the second quotation, that removing the mutual right to kill would render war more chaotic and immoral. Why should we think that a rule allowing mutual killing of victims will be more successful in maintaining order and morality than one forbidding, or at least limiting, such killing? One (I think very compelling) answer is that such a rule is far more likely to be followed, and I discuss this further when I present my own position below. However, it is not the sort of answer Walzer wants to give. As I discuss below, Walzer wants to show that a soldier's right to kill is a robust moral right which no-one may take from him, and that when a soldier kills, what he does is truly blameless, truly morally inoffensive, rather than merely excusable. But if the only reason soldiers may kill one another is that, as it happens, without this permission, war would become much worse, Walzer's 'right' to kill would be reduced to something far less robust: not a right at all, but a contingent freedom dependent entirely on utilitarian considerations. (In fact, his position would be reconciled with my own.) Therefore I think further argument is needed before Walzer's position is rendered appealing.

The importance of consent: excuse and justification

The fact that soldiers do not consent to their situation does in fact affect our moral assessment of their behaviour, and by appealing to the victim-status of soldiers, Walzer might make a credible case for the claim that soldiers are not to blame for the acts of war they perpetrate. But there are two ways soldiers might be freed of blame for killing: a soldier's act of killing may be an *excused* act, that is, an act for which the agent is exempt of responsibility by reason of duress, or madness, or some such consideration. Or it may be a *justified* act, an act which one is not blameworthy for performing.¹⁴ Walzer seems to want to endorse the latter account of a soldier's lack of blame for killing, but I shall argue, with Rodin, that his position is more consistent with the former.

¹⁴ For further discussion of the distinction between excuse and justification, see Rodin, pp.90-96, pp.170-173. Rodin discusses the analogy between the standard of duress appealed to by a soldier, and that accepted in common law. Another type of problem for Walzer might be that duress is not accepted as an excuse for homicide in common law (Rodin, p. 171); should it be accepted as one in war? If not, then even if soldiers do not consent to their position, they still are not excused from their acts of killing. Thus excuse and justification may both be absent; and any grounding for a mutual entitlement to kill goes with them. (It is also worth noting here that the language of excuse and justification has no place in utilitarian moral theory, according to which any permissible act will be justified.)

To illustrate the difference between an excused and a justified act, consider this example: If I am forced at gunpoint to steal something, I may be *excused* from this act of stealing, but it does not follow that I had or have a *right* to steal. Likewise, while soldiers may be *excused* from responsibility for killing, it does not follow from this alone that they possess a *right* to kill. In fact, in certain circumstances where the excuse was insufficiently strong (say if he joined the army freely knowing he would be fighting for an unjust cause) a soldier might no longer be excused from his action – just as I, if I knew the gunman’s gun was fake and he presented no real threat to my safety, would no longer be excused from stealing at his demand. We can gain and lose these sorts of entitlements to perform usually prohibited actions, because we only have them when the excuse is sufficient. But I do not think that Walzer wants to say that a soldier’s act of killing is a bad action which is excused by reasons of duress, allowed by this sort of provisional entitlement. Rather, he wants to argue that a soldier’s act of killing it is not a bad action at all, but a fully *justified* one.

Here we run into a circularity problem, identified and discussed by Rodin on p. 172 of *War and Self-Defense*. According to Walzer, soldiers are not in need of excuse when they kill one another, because they do not act badly - they are merely exercising their mutual entitlement to kill. This entitlement derives from the fact that they are mutually non-consenting victims. However, as we have seen, the fact that soldiers do not consent to their situation serves not to justify, but to *excuse* their actions. Thus Walzer claims that a soldier’s act of killing is justified simply because it is excused, a false inference. Walzer needs to provide another reason for thinking that a soldier’s act of killing is not merely excused, but justified – and that justification cannot make reference to his grounds for the mutual entitlement to kill. Rodin sums up the problem neatly here:

[Walzer] shifts from claiming that soldiers are excused for killing because their lack of consent means they are not responsible, to the claim that their killing is justified because they possess an entitlement to kill . . . [But] when we ask where the mutual entitlement to kill could

derive from, it looks very much as if it could only arise out of the kind of mutual consent which Walzer has already denied in respect of soldiers.¹⁵

I find Walzer's duality of reasoning to be at times very obvious. Consider this quotation from chapter 3 of *Just and Unjust Wars*, in a section entitled 'The Moral Equality of Soldiers':

When soldiers fight freely, choosing one another as enemies and designing their own battles, their war is not a crime; when they fight without freedom, their war is not their crime.¹⁶

What Walzer is saying here is that when there is mutual consent between soldiers, their acts of killing are justified, non-criminal acts, 'not a crime'. When there is mutual lack of consent, their acts of killing are excused acts, 'not their crime' – but the acts are, in these circumstances, still criminal acts of killing. If all soldiers really possessed a mutual right to kill, in the sense that their actions were justified and not merely excused, their acts of killing would not constitute a crime – and contrary to what Walzer suggests, no blame for them could rest with the statesmen who brought them about.¹⁷

I think Walzer's confusion of excuse and justification in grounding the mutual entitlement to kill can be made most apparent when we present him with this situation: Imagine a soldier freely consents to join an army and engage in battle. He knows that his enemy is fighting without consent. From what derives his right to kill that enemy? He is not *excused* from his act of killing, because he acted with consent; and yet his killing is not *justified*, for his opponent acted without consent, and as Walzer recognises, war is only 'not a crime' when both parties act freely. In this case the soldier's right to kill seems unrelated to any notion of consent. Is there another kind of justification for it?

¹⁵ Rodin, p. 172.

¹⁶ Walzer, p. 37

¹⁷ See Walzer, pp.34-40. Walzer's shifting of responsibility for the evils of war from soldiers to politicians is explicitly stated on more than one occasion. Again, note that if soldiers possessed a mutual entitlement to kill, and their actions were justified, then the killing of soldiers could not be counted as the 'evils of war' for which politicians are guilty. But Walzer evidently thinks the deaths of soldiers ought to be counted among these evils; see p.29ff, 'The Tyranny of War'.

A pragmatic justification for the mutual entitlement to kill

I have argued that we should reject the notion that a mutual entitlement to kill derives either from the consent of the soldiers involved, or from a lack of such consent. As I stated above, Rodin's discussion up to this point follows a similar course to mine, but our conclusions now diverge. Rodin, unable to identify any other feature of soldiers representing a promising avenue for grounding the entitlement, concludes that 'it would seem that a long-standing tenet of Just War Theory must be abandoned – soldiers fighting an unjust war have no permission to kill, and there is no 'moral equality' between soldiers.'¹⁸ Like Rodin, I think there is no intrinsic feature common to soldiers on both sides of a war which grounds their mutual entitlement to kill. I agree that 'soldiers fighting an unjust war have no permission to kill', if that permission must stem from an intrinsic entitlement or right. But I believe that there is still a utilitarian reason to prefer rules of *jus in bello* which ignore this moral asymmetry between soldiers. To remind ourselves of our original puzzle: it is strongly intuitive that soldiers fighting for an unjust side ought not to be allowed to target soldiers fighting for a just one. So why do we think that the 'entitlements and restrictions' of *jus in bello* apply to both the unjust and the just (if there is one) side in war? How can we justify rejecting the intuitively obvious in preference to the intuitively problematic?

My answer is this: first, it is not problematic or paradoxical to accept that there is no intrinsic or deontological moral property (such as consent, duress, innocence, etc.) common to all parties concerned which can explain why an army fighting an unjust war is justified in attacking the soldiers of its just enemy. This fact merely indicates that if justification for the entitlement is to be found, it must be found by an appeal to different sorts of moral considerations – and as I stated above, I believe that utilitarian considerations can justify the mutual entitlement to kill. The utilitarian can argue that it is generative of the best outcomes to accord both sides equal permission to kill one another, or rather to formulate rules of warfare which will achieve the best results when followed by both sides of a dispute.¹⁹ This is because it would be futile to

¹⁸ Rodin, p. 173

¹⁹ This may sound like I am forcing a utilitarian justification on the Just War theorist at the outset, and I would dearly love to do so, but in fact I think any kind of Just War theorist can (and ought) endorse the

restrict permission to kill in war to one side of the conflict, given that in real-world situations the members of the ‘unjust army’ are neither willing nor obliged to identify themselves as such. When states go to war, they invariably declare (whether through actual conviction or motivational convenience) that they are fighting a just war, and that their enemy is not. *Jus in bello* must therefore take account of the fact that in real-world warring situations, anybody needing to apply its rules will claim to have already satisfied the requirements of *jus ad bellum*. Any Just War theorist who wants to make the actual fighting of wars less bloody must generate her principles on the assumption that they will be followed by *both* warring parties.²⁰

This is not to say that any side adhering to the principles outlined in *jus in bello* will be acting in an irreproachable manner in utilitarian terms, for although the wrongness of fighting for an unjust cause may be increased by fighting unjustly for it, it cannot be erased by fighting justly for it.²¹ However, given that our ability to alter the behaviour of those fighting for an unjust cause is so limited, there is good reason to formulate rules of warfare which will generate good consequences within these limitations. However we might judge the isolated actions of soldiers or of heads of state, whether as utilitarians or as rights theorists, these judgements need not be perfectly reflected in those rules which we think ought to be legislated for the conduct of war. This understanding of the difference between ideal moral judgement of individual acts, and decisions about which sorts of rules or laws ought to be in place in a non-ideal world, explains the apparent paradox identified and discussed by

minimal consequentialist standard of requiring that one’s moral edicts actually have some chance of being followed – a perverse application of ‘ought implies can’, if you will. Even a Kantian can agree that it is not good Just War theory solely to state that ‘it is wrong for those fighting an unjust war to kill the enemy’, no matter how true his general moral theory might show this to be. The nature of the Just War project requires an approach which is influenced, to at least some degree, by the likelihood of the course of action one recommends being followed. Where different moral theories will diverge is in judging which sorts of feasible courses of action are preferable.

²⁰ I discuss the conflict between ideal morality and political reality further, and along similar lines to those above, in section 3 of chapter 5.

²¹ It is difficult, in discussing utilitarian and deontological moral concerns together, to avoid language which is incompatible with at least one position; thus I use the terms ‘justice’ and ‘rights’ to mean what they mean in the Just War tradition. ‘Just wars’ are, for my purposes, those entered into according to the standards of *jus ad bellum* and fought according to *jus in bello*, without reference to what the utilitarian will say of the rightness of the overall endeavour.

Having said this, I think it reasonable to suppose that wars fought for ‘unjust’ causes (as a deontologist will identify them) are likely to be more generative of harm than those fought for ‘just’ causes. That is, aggressive wars fought for gain are likely to be more harmful, at least in the long term, than wars fought to liberate oppressed persons or to defend a valuable way of life. Therefore, even in utilitarian terms, there is good *prima facie* reason to prefer ‘just’ warfare to ‘unjust’.

Rodin. Rodin's concern was that when an aggressive army adheres to *jus in bello*, 'the war taken as a whole is a crime, [but] each of the individual acts which together constitute the aggressive war are entirely lawful'. Now we can see that the acts of war being *lawful*, in the sense that they comply with *jus in bello*, need not mean that they can be justified by appeal to intrinsic facts about the parties involved.²² Rather, the law which permits them is justified; and it is justified by appeal to utilitarian considerations.²³ This difference between what is legally appropriate for the conduct of war, and what is morally justified in moral isolation or under ideal conditions, arises because there is no sufficient law-enforcement which would allow very stringent laws to be formulated and enforced to good effect. Thus the laws of *jus in bello* must be designed to operate to best effect in a kind of artificial (and impossible) moral universe in which we assume that all sides fight with a just cause.^{24, 25}

The direction of my arguments

The arguments I give throughout the following discussion of non-combatant immunity are designed to apply to both warring parties regardless of the justice of their cause. The task is to identify what freedoms and constraints ought to apply when the opposing army is similarly free and constrained in their actions; and I argue (with particular reference to non-combatant immunity) that a utilitarian conception of the rules of warfare is the only morally defensible one.

As will become clear, my account of what ought to be permissible under *jus in bello* is often more restrictive than the Just War tradition dictates, but it is also potentially

²² In legal terms, there is no suggestion of paradox. The legal claim would be that an army adhering to *jus in bello* commits no crimes, and yet their war, if it contravenes the requirements of *jus ad bellum*, is illegal. This is unproblematic because in law, it is easily possible for an overall endeavour to be illegal, even when the individual acts which comprise it are usually not. For example, it is not usually illegal for me to drive my car to the bank, pick you up, and drop you off somewhere later; but it *is* illegal for me to do this if I know you have robbed the bank and my car is the getaway car – it is illegal even if I adhere to all the relevant traffic laws. Likewise with soldiers justly fighting an unjust war: it is possible that their overall endeavour be illegal even if they adhere to all the rules of warfare.

²³ This is not a rule utilitarian account of *jus in bello*. As I discuss further in later chapters, it will always be the case that, when the rule of law can be flouted in order to generate better consequences, this should be done. The law is only justified insofar as it *actually will* generate the best consequences.

²⁴ Once again, this 'artificial moral universe' is only impossible for the deontologist. Utilitarians require only that we endeavour to bring about the best possible consequences, and sometimes these consequences will include 'unjust' acts.

²⁵ One might think that it would be ideologically preferable for *jus in bello* to acknowledge that soldiers fighting an unjust war do *not* have a right to kill the enemy, but the practical effects of this doctrinal change would be negligible. At most, it would enable third parties to the war to condemn those fighting an unjust war with reference to *jus in bello* as well as *jus ad bellum*.

more permissive. The account is potentially permissive because what counts as a legitimate act of war will vary in light of the ‘last resort’ criterion of *jus ad bellum*. Loosely, some ends (the prevention of global nuclear warfare, for example) seem to be justifiable at almost any cost, and if usually forbidden acts of war are the best or only way to attain them, then these acts may be rendered morally acceptable, or even obligatory, as means to this desirable end.²⁶ On the other hand, the utilitarian nature of my arguments also leads me to conclude that in most cases, justifiable wars and justifiable acts therein are more limited than we usually imagine, because we are obliged to adopt the *least harmful means possible* to our ends. Where feasible this will require us to avoid resort to war all together. Where it is not possible to avoid war, this obligation will generate a strong requirement to demonstrate the military necessity of all deaths, both civilian and soldierly.

The remainder of this thesis is devoted, first, to explaining why a utilitarian ethic is required to salvage the criterion of non-combatant immunity which lies at the heart of just war theory; and second, to explaining how, using this type of ethic, we can correctly identify permissible acts of war.

²⁶ This kind of permissiveness would rely on reliable international law-enforcement to ensure the rules were not abused. I discuss the importance of a reliable body of international law in chapter 5.

CHAPTER TWO

Innocence and guilt

In section 2 of the previous chapter I introduced the idea that we often appeal to considerations of the innocence or guilt of an agent when we wish to justify killing. The dual notions of combatant guilt and civilian innocence have been appealed to throughout the history of immunity theory to explain why it is that killing civilians is prohibited while killing soldiers is acceptable. In this chapter I examine the strength of these types of explanations. First I question what is meant by the ‘innocence’ of civilians, and the ‘guilt’ of combatants, and give some reasons to doubt the appropriateness of these labels. I also question the assumption that these considerations permit us to take or protect life with the stringency supposed by immunity theorists. I then argue that even if we accept that guilt allows us to kill soldiers in wartime, and that innocence prevents us from killing civilians, still these principles will not endorse anything resembling immunity theory as we know it. I conclude that either we must accept that innocence and guilt are not important in supporting the principle of non-combatant immunity, or we must agree to endorse a model of non-combatant immunity which differs significantly from that found in the Just War tradition.

Section 1: A definition of innocence

What it means to call a non-combatant ‘innocent’ in the context of *jus in bello* stands in need of clarification. One analysis of Just War theory considers ‘innocence’ simply to mean ‘not posing a threat’. David Rodin, in *War and Self-Defense*, clearly accepts and endorses this characterisation:

The non-innocent (those who may permissibly be killed in war) are not defined as morally non-innocent or at fault, but rather as something like ‘presently harmful’.²⁷

He thinks such a characterisation is necessary if we are to reconcile three statements of Just War theory which seem to be mutually incompatible:

It is only through such an objective (re)definition of innocence and non-innocence that the Just War Theory is able to maintain what would otherwise appear to be an inconsistent triad of propositions: (i) soldiers on either side of a war who abide by the rules of *jus in bello*, or the war convention, are morally innocent, (ii) one is never permitted to kill the innocent, (iii) one is permitted to kill enemy soldiers in war.²⁸

I think that Rodin is right to say that a definition of innocence which refers to the intrinsic moral status of soldiers and civilians is problematic, but not for the reasons of inconsistency that he identifies here. Indeed, it is not clear that a consistent redefinition of innocence such as Rodin describes would help us reconcile these three propositions, because statement (i) would then make the false claim that soldiers fighting in accordance with *jus in bello* are not presently harmful. If we are to reconcile the propositions by manipulating the use of the word ‘innocent’, then we must equivocate between the two interpretations of that word, such that in (i) it retains its traditional sense, and in (ii) it means ‘not presently harmful’. But I see no reason to think that (ii) should properly be read as ‘one is never permitted to kill the not-presently-harmful’: this perverts what is usually taken as a straightforward appeal to an intrinsic moral property. Moreover, if the Just War theorist does accept this reading, she still faces a potentially serious problem: *jus in bello* will now state that we may sometimes kill the morally innocent, and the justification for this will be that the morally innocent sometimes pose a threat. The natural conclusion of this line of argument is that what justifies targeting persons in war, whether they are morally innocent or not, is the level of threat they pose. As I discuss further in chapter 4, this conclusion will spell bad news for the immunity theorist, because no stipulated level

²⁷ Rodin, p. 84

²⁸ Rodin, p. 84, footnote.

of threat will serve to discriminate between soldiers and civilians in the way non-combatant immunity requires.

In summary, if the Just War theorist supports an equivocal definition of innocence to allow *jus in bello* to retain its internal consistency, she does so at the expense of non-combatant immunity, one of that doctrine's central tenets. She can avoid this conclusion either by accepting that Just War theory really is logically inconsistent in the way suggested by Rodin, not an attractive avenue; or by claiming that (at least) one of the three statements he identifies as characterising it is false. I have explained in my discussion of the separation of *jus ad bellum* from *jus in bello* (section 3 of chapter 1), I think there is very good reason for rejecting (i).²⁹ Alternatively, the Just War theorist might take issue with the absolute nature of the moral prohibition contained in (ii): after all, other parts of Just War theory, such as the doctrine of double effect, seems to allow exceptions to the ban on killing innocents.

So, *pace* Rodin, the Just War theorist cannot appeal to an unequivocal redefinition of innocence to reconcile her theory's tenets. Even if such a redefinition could be made to work, I think that there are good reasons for rejecting this characterisation of non-combatant innocence. First, it renders the common cross-cut definition of a non-combatant, as someone who is *both* innocent *and* who poses no threat to the enemy, redundant.³⁰ If moral innocence is not part of the characterisation of non-combatants, then to refer to their innocence as offering them protection from attack does not have the clear intuitive appeal it seems, but is rather a foreign and misleading use of a common moral term. When the immunity theorist says that civilians ought not to be attacked because they are innocent, it is the notion of *moral* innocence, not non-harmfulness, which gives her claim such intuitive appeal. (I think the immunity theorist would do well to value this kind of moral innocence; but as I shall argue below, she cannot successfully use it to distinguish between combatants and non-combatants.) But suppose that the immunity theorist is properly understood as claiming that civilians ought not to be targeted just because they do not pose a threat

²⁹ In fact, Rodin himself rejects this statement, arguing that soldiers may be held personally responsible for fighting in an unjust war. See Rodin, p. 165 ff, 'The Responsibility of Soldiers.' It is interesting that he does not revisit his mention of the definition of 'innocence' as a result, as he no longer is committed to this redefinition in order to avoid inconsistency.

³⁰ See section 3 of chapter 4 for further discussion of this cross-cut definition.

to the enemy. In this case, not only does she face the argument that this position is untenable (as I argue in chapter 4), but she faces the additional objection that this claim is not what we agree to when we agree that innocence ought to provide moral protection.

This is particularly clear in the case of children: when we agree that children ought not to be targeted in war because they are innocent, it is not their lack of threat to which we refer, but their lack of responsibility for their situation. Even children who do pose a threat to the enemy, such as child-soldiers or children forced to work in munitions factories, arguably retain their moral innocence in the sense relevant to whether they deserve protection from killing. The constitution of a threat does not (or need not) equate to a lack of moral innocence, and when we consider whether ‘innocence’ ought to offer protection, the agent’s *moral* innocence is more important than whether he constitutes a threat.³¹

In summary, it seems to me reasonable to think that when we say non-combatants are innocent and ought not to be targeted, we do not merely mean that they do not pose a threat. Rather, we mean this to be something like a statement about their lack of moral responsibility for their involvement in the war.³² It is this notion of innocence which I discuss below.

³¹ This is, in fact, a very basic problem for immunity theory: there seems no good reason to suppose that moral innocence and the constitution of a threat are logically connected in any way. One is a feature of an agent’s causal behaviour, and the other, of her intentional behaviour. It is inappropriate to pass judgement on an agent’s culpability purely by looking at the consequences of her behaviour. Likewise, it would be not only inappropriate, but actually impossible to decide whether someone poses a threat just by considering whether she is responsible for her actions. I return to this distinction in greater detail when I discuss innocent aggressors and self defence, in chapter 3.

³² Of course, the immunity theorist might reply that although we do not generally mean this, and although our intuitions may not be answered by this use of the term ‘innocent’, she can nevertheless make a case for her position independent of these intuitions. To this I would say two things: First, the immunity theorist has a lot of uphill work to do if she is reduced to arguing that non-combatant immunity can be justified solely by an appeal to the threat posed by soldiers. This is not just because I think such arguments fail (see chapters 3 and 4), but because so many supporters of non-combatant immunity rely on the widely-accepted notion of ‘innocent civilians’ to motivate their position. If you tell an immunity-theorist-sympathiser that what makes it wrong for the military to target schoolchildren is not that they are morally innocent, but rather that these children don’t pose any threat, the response is unlikely to be hearty agreement.

It is true that this argument is only about our moral intuitions, which are sometimes inappropriate. The immunity theorist can certainly bite the bullet and say that although her position may be counter-intuitive, it is also the only one capable of logical defence. But I feel certain that she does not want to say this. Rather, I suspect she wants both to harness our intuitions about why civilian life is precious, and reject our reasons for thinking so; have her heuristic cake and eat it too.

Section 2: Deciding the guilt of soldiers and the innocence of civilians

A distinction between combatants and non-combatants based on an appeal to innocence is attractive because innocence is a moral concept which we widely accept as important.³³ We point to issues of innocence and guilt when justifying the taking of life, usually with reference to a system of law under which death is a punishment for some types of crime. Proponents of the death penalty will generally hold that the criminal's guilt is of primary importance in justifying depriving him of his life. (Pro-death-penalty arguments are controversial, but I will assume here that some actions are justly punishable by death. If not, then no matter how 'guilty' soldiers were decided to be, we could never use this standard to distinguish them from civilians when it comes to taking life.) If civilians are innocent in a sense in which soldiers are not, then a principle allowing attacks on the one group and not the other will, if appropriately applied, have moral weight. The targeting of combatants could be seen as some kind of 'punishment' for which civilians are not liable due to their non-involvement in the 'crimes' of war.³⁴

Many people do think that civilians are innocent in a sense in which soldiers are not, and the term 'innocents' is often used interchangeably with 'civilians' and 'non-combatants' in Just War literature. It is easy to see why this idea is so widespread: civilians are not directly responsible for killing anybody, and in many non-Western countries the rural communities in particular know little of the political situation, do

³³ This point is discussed by George Mavrodes ('Conventions and the Morality of War', *Philosophy and Public Affairs*, vol.4, no. 2, 1975, p.121), whose arguments I turn to in greater detail in chapter 5. Also, note that I use the terms 'guilt' and 'non-innocence' interchangeably throughout these discussions.

³⁴ Rodin discusses at length some of the problems we face when we try to treat the combatant as a criminal or to view war as law-enforcement: 'There is a gap in the moral explanation between a right to act against an aggressive state, and the right to act against the persons who are its soldiers – a conceptual lacuna between the two levels of war' (p. 165). Also: 'Surely if it is justifiable to inflict harm against the soldiers of an aggressive state, this is because they share responsibility for the wrongfulness of the aggressive war. What is curious is that the Just War Theory precludes appeal to such considerations since it denies that ordinary soldiers are personally responsible for fighting in an unjust war' (p. 166, general discussion pp.165-180).

Although I think there are real problems for the immunity theorist to face here, I will not discuss these issues. This is because I think that arguments from innocence fail even if it *were* appropriate to think of combatants as criminals and to look upon war as an exercise in punishment. Even when given the maximum degree of charity, and maximum benefit of any doubt surrounding auxiliary claims, arguments from innocence still cannot endorse non-combatant immunity.

not welcome fighting, and have done nothing to advance the war. Even civilians working in factories supplying the military are often employed in ways which exist independent of warfare; food and clothing are not uniquely military requirements. In the case of children, it is especially easy to see that they are in no way responsible for the situation in which their country lies. We usually accept that this non-responsibility extends to much of the general civilian population.³⁵

If the immunity theorist thinks that it is civilians' non-involvement in war and their non-responsibility for it which renders them innocent and worthy of protection, then she must think that it is their direct involvement in acts of war and their responsibility for this involvement which renders combatants guilty and worthy of attack. In the remainder of this chapter I carefully consider the basis for these judgements and the difficulty in applying them. If combatants are to provide an exception to the moral wrongness of killing on the grounds of their guilt, they must both perform guilty actions, and be responsible for those actions. Similarly, if civilians are to be protected by their innocence, then they must both refrain from performing guilty actions, and refrain from being morally responsible for the guilty actions of others. Below I argue that neither of these standards is consistently met and thus arguments from innocence, despite their intuitive appeal, fail to support the principle of non-combatant immunity.

Three problem cases for arguments from innocence

If we accept that civilians are innocent because of their lack of connection with, or responsibility for, war, then I argue that we must also accept the following three classes of soldiers as likewise innocent:

1. Soldiers who are drafted into the army
2. Soldiers who joined the army as a result of propaganda or misinformation

³⁵ Contrary to this, Igor Primoratz has argued that responsibility for war ought to be extended to all voting citizens living in democracies (Primoratz, I., 'Michael Walzer's Just War Theory: Some Issues of Responsibility'. In *Ethical Theory and Moral Practice*, vol. 5, no. 2, June 2002, pp. 221-243). He argues that even those citizens who voted *against* the government which chooses to go to war bear some responsibility for that choice, and this responsibility can only be removed by their opposing the war to the best of their ability. While I find these arguments interesting, I think the notion of responsibility identified by Primoratz is probably insufficient to justify attacks on civilians. There is also an important question about how effective opposition to the war will be: there will be cases in which it cannot be expected to make any difference, and it seems strange to consider an exercise in futility obligatory.

3. Soldiers who are justly fighting a just war

I argue that soldiers in these three groups are not morally responsible, or at fault, for the circumstances in which they become soldiers, and thus the fact that they are soldiers can tell us nothing about their moral innocence or guilt. The remainder of this chapter is an examination of these three classes and of some objections that might be raised against my arguments.

Innocent draftees

In countries where soldiers are drafted into the army, the distinction between the innocent civilian and the civilian who becomes a soldier and thus loses his innocence will depend on age, sex and general health. We do not usually consider these to be moral categories. If we are to hold the draftee responsible for his actions, then, it seems we must argue that he had some control over his behaviour with respect to the war: In other words, we must hold that he consents to his position in a strong enough sense that he is culpable for actions performed as a result of being in that position.

In the Western world, this is probably a fair judgement. It is reasonable to think that that if a Western soldier truly did not wish to be part of the army, he need not have been, and therefore he must accept at least *some* responsibility for his situation.³⁶ But in many countries, refusing the draft is legally punishable, even by death. In these cases it would be wrong to say that a drafted soldier had the real freedom of choice

³⁶ Whether this responsibility is strong enough to justify killing him is questionable. After all, if he were perhaps a year or two younger, or female rather than male, he would have retained his innocence with no effort on his part. Is his age or sex, coupled with his choice not to refuse the draft, sufficient to condemn him to death? It may be useful to consider an analogy from civilian law: Imagine that my local government requires me, if I do not wish to pay a certain special tax, to declare my opposition to the tax on the first Sunday of the month in the town square at 9am. This requirement applies only to a certain group of men in my state – the others aren't asked to pay the tax at all and need do nothing to avoid it. If I refrain from making the required announcement, am I then morally obliged to pay the tax? Possibly to some degree; after all, the burden of announcing my opposition is not a heavy one. However, it does seem unfair that I should have to say anything at all, when others need not. Even if I am to some extent obliged to pay the tax, it seems wrong to punish me if I do not (let alone to punish me by death).

An important difference between this example and the example of the soldier is that while my refusal to object against the special tax will result only in my paying a small amount to the local government, my failure to refuse the draft may lead me to kill large numbers of innocent people. We might say that when a failure to object against some rule leads us to perform these kinds of actions, the obligation to object is much stronger than when it leads to morally insignificant actions. Still, remember that the soldier who fails to refuse the draft is 'punished' by becoming vulnerable to just attack, effectively losing his 'right to life'. This seems unduly harsh. Whatever his 'crime' in accepting the draft, it does not seem it ought to be punishable by death.

required for consent. If he is forced to stop being a civilian against his will and upon pain of death, there seems no reason to believe he loses any of that innocence which, for example, a twin sister of his would retain. If this is the case, then any army into which soldiers are drafted, and in which any coercive measures are used to ensure the draft is not refused, will contain at least some soldiers who have done nothing beyond being young, fit and male to lose their moral innocence, and the protection against killing it affords.³⁷

An addition to this category of innocent draftees is that of child soldiers. In Liberia, Burma, Sierra Leone, Colombia, and many other countries troubled by civil conflict, 12-year-old soldiers are involved in acts of war to the same extent as their adult counterparts. But despite their physical involvement in ‘criminal’ acts, the moral innocence of these children is unquestionable; and if moral innocence offers any protection from enemy attack, these children ought to be protected.³⁸

Innocent victims of propaganda and misinformation

The moral relevance of coercion to responsibility for one’s actions leads us to the second category of innocent soldiers. These are those soldiers who join in the war effort as a result of misinformation. As an example, consider an isolated rural community which is subjected to government-supplied non-factual information about a probable future war and its likely outcomes, its aims, the justice of its cause, etc. If, as a result of this misinformation, some young men of that community decide to join up as soldiers, then those young men lose their moral innocence only if they would

³⁷ It is true that in many countries, refusing the draft will not lead to anything as serious as execution. However, it will often be punishable by a prison sentence or fine. Even a relatively mild punishment-based incentive is sufficient to illustrate the moral point: Suppose that war is being waged by Johnny’s country and soldiers are being drafted. Johnny, though not generally a pacifist, believes that the war he is being asked to fight is not a just one. However, refusing the draft results in a 3 year prison sentence, and wishing to avoid this, Johnny joins the army. In such a case we may well think that Johnny has done something wrong. After all, if his principles really lead him to believe the war is unjust, he ought not to be persuaded by a limited prison sentence to act immorally. However, the wrong Johnny has done does not seem serious enough to make him a morally justifiable target for death. Imagine it as a criminal offence: Johnny is convicted of acting immorally because he chose to become a soldier even though he knew he would be fighting an unjust war. Would the death penalty be an acceptable penalty for such a crime? If not, then why does this type of coercion into becoming a soldier mean that that soldier loses that innocence which a civilian retains, the innocence which means he may not be killed?

³⁸ Comprehensive information about child soldiery worldwide can be found at www.child-soldiers.org. The important distinction between posing a threat and being culpable is particularly sharp in the case of child soldiers, and our intuitions about when we may kill those who threaten us are particularly confused. I consider these issues further in chapter 3, when I discuss our permission to defend ourselves against innocent threats.

have joined up anyway had the information on which they acted been factual, or if their belief in the misinformation was unreasonable. After all, a decision made on the basis of reasonably believed but false information, such as might be supplied by a manipulative source, is not a decision for which the agent concerned is responsible. The proper conditions for free and responsible decision-making are not present.³⁹

It is true that it is dangerous to allow propaganda to provide too much absolution from responsibility. We do not want to say, for example, that Nazi government officials were not responsible for their actions because they were given misinformation about the moral status of the Jewish people. But then, it is not the case that Nazi government officials received this information in a moral void or were unaware of the motivation behind its dissemination. There was both foreign and domestic political dissent against the information propagated, and in addition the anti-Semitic claims made were contrary to what was already (professedly) widely believed about the equality of persons. When a claim with important moral implications is the subject of this kind of dispute, an individual has a responsibility to assess the arguments for and against its acceptance before making any decision which assumes its truth or falsity. However, where there exists only one set of information, which is presented as fact by those we reasonably expect to be informed, then attempting or insisting upon independent verification of that information is often irrational, or even impossible. For example, when scientists tell me that the planet Earth is an oblate spheroid, I accept it as true. If I were to refuse their authority and instead insist upon finding out for myself, I would be expending futile energy. The authority of the scientific community is given a certain amount of credibility and it is reasonable of me to respect it.⁴⁰ I would suggest that the government of one's own country, when no dissenting voices are heard or such voices are silenced before they are heard, is given a similar degree of credibility – and those who believe what they are told in such circumstances are not negligent in doing so.

³⁹ Having said this, sometimes an agent will choose not to avail himself of relevant, freely attainable information about his decision, in which case he *is* responsible for his decision, and may also be culpably negligent. An example might be a parent who diagnoses his child's illness without referring to recorded medical knowledge.

⁴⁰ It is important to note that this is not because the information is undisputed. There are certainly people who deny that the Earth is an oblate spheroid. Various manifestations of the Flat Earth society can be found at: www.flat-earth.org/society/about.html, www.lhup.edu/~dsimanek/fe-scidi.htm, and www.alaska.net/~clund/e_djublonskopf/Flatearthsociety.htm, amongst many others. But these people do not have the kind of credibility that the scientific community has, and rightly so.

Because the Western world has wide access to alternative media and reasonably wide access to free speech, because dissent is voiced to nearly all significant governmental decisions, and most importantly because of our high rate of literacy, we may tend to think that being taken in by propaganda is an example of gullibility or wilful self-deceit. If we doubt the representation offered by our government of a given situation, we can find out for ourselves what the fact of the matter is. One might object that the picture I suggest of isolated communities without anti-governmental tendencies is a naïve or even an offensive one. But I do not consider it unrealistic, especially where written material is unable to be employed to offer an alternative viewpoint. Sometimes governments can access information which their citizens cannot, and in such cases the citizens must rely on the reports of their government. If the government of a nation whose rural citizens were without access to telecommunications were to tell those citizens that a neighbouring state was preparing to go to war with them, we would not require that each citizen independently verified this information before acting on it.

A second avenue of support for my claims comes from the historical consideration of other non-majority movements and how their claims were initially received. Groups first speaking out against nuclear power, for example, or in favour of black civil rights or votes for women, were publicly ridiculed, socially marginalised, and seen as fools, trouble-makers, or criminally insane. People simply do not always, or even usually, respond with open and interested minds to alternative points of view, and especially not when one opinion is given by an official body and the other by a small group of 'social misfits'. Even though we now think that those people who derided the suffragettes were making their decisions based on wrong information, we also accept a huge history of mitigating factors in their favour. We certainly do not think that they ought to pay with their lives for their moral errors. Why so, then, in the case of soldiers entering into armed service under the weight of a similar history of misinformation?

It is important to recognise that only a very small class of soldiers will be excused from their ignorance about the justice of the war in which they are engaged, and this class is probably growing smaller every day as a result of the spread of internet

technology. However, I say again: Any soldier who non-culpably but falsely believes that he ought to fight for his country has done nothing to make him non-innocent in the way necessary for him to be justly be targeted and killed by the enemy.

Innocent soldiers justly fighting a just war

I have discussed in section 3 of chapter 1 why I think it is important that rules of warfare be designed to be followed by all parties to the conflict on the artificial assumption that they are all just in their cause. But as well as this pragmatic requirement, it is also important to acknowledge the possibility of some party to a conflict actually justly fighting a war with just cause according to the specifications of the Just War doctrine. If we grant that this is possible, then there seems no justification for the claim that the soldiers of such an army have done anything to forfeit the moral protection their innocence would give them. Indeed, they seem rather to be more morally worthy than the civilian who refuses to become involved in the war, and the more responsible they are for their actions (i.e. the less coercion used to encourage them to become soldiers, the more accurate the information on which they acted, and the more fulsome their consent), the more worthy they will be.

Of course, in real world warfare no government is going to acknowledge that their enemy is justly fighting a just war, nor is any individual soldier very likely to make the same admission. Without some international judicial body with sufficient authority to make such a ruling, this class of ‘innocent soldiers’ will never be identifiable. For the argument above to be undermined, however, these soldiers would have to be not just unidentifiable, but nonexistent; and there is no reason to suppose this will always be the case.

The problem of politicians

Politicians and other leaders represent a fourth, slightly different problem case for any argument for non-combatant immunity based on a concept of moral innocence. Instead of a class of soldiers who seem innocent in respect of war, they are a class of civilians who seem culpably responsible. George Mavrodes, in ‘Conventions and the Morality of War’, asks us to consider someone who is ‘an enthusiastic supporter of the unjust war and its unjust aims’:

He may give to it his voice and his vote, he may have done everything in his power to procure it when it was yet but a prospect, now that it is in progress he may contribute to his both his savings and the work which he knows best how to do, and he may avidly hope to share in the unjust gains which will follow if the war is successful.⁴¹

Such a person is, assuming the work they know best how to do is non-military, clearly a civilian. Even a politician who provokes or declares war is, unless he holds some military post, considered a civilian. Mavrodes compares these cases to ‘a young man of limited mental ability and almost no education’ who is drafted and sent to the front with no comprehension of the meaning of the war and no particular wish to be engaged in it.⁴² In terms of the combatant/non-combatant distinction, the young man is clearly engaged in warfare and is therefore non-innocent. The supporter of the war, conversely, is clearly not engaged in warfare and so is innocent. But this application of the terms ‘innocent’ and ‘guilty’ seems to have no basis in morality whatever. As Mavrodes puts it: ‘We must see to it, if we are to be moral, that we punish someone only for his own crime and not for someone else’s.’⁴³ If we punish the drafted soldier by making him a legitimate target, but not the politician or business leader who placed him where he is, we are guilty of exactly that miscarriage of justice.

Section 3: Justifying the killing of innocent soldiers

The three cases of ‘innocent soldiers’ outlined above give us good reasons to think that at least some soldiers are innocent in the same way that most civilians are.⁴⁴ This poses a problem for those who wish to use the idea of innocence to defend the protection granted to non-combatants: If civilians are to be immune from attack because of their innocence, it seems as though these innocent soldiers ought to be as well. There are two responses an immunity theorist might give to this problem. The first is to say that because at least some of the soldiers, and probably the majority, are

⁴¹ Mavrodes, G., ‘Conventions and the Morality of War’ (*Philosophy and Public Affairs* 1975: vol. 4 no. 2, pp. 117-131), p. 122.

⁴² Mavrodes, p. 123

⁴³ Mavrodes, p. 123

⁴⁴ The actual numbers will vary from conflict to conflict, as there is huge variation between nations concerning the application of the draft and the penalty attached for refusal, and the prevalence of propaganda and coercive measures to pressure people to enlist.

not innocent, then soldiers as a group are rendered legitimate targets. I discuss this idea further below. The second option is to launch a pragmatic argument concerning the impossibility of only targeting non-innocent soldiers. This, I claim, is either to desert the ‘intrinsic difference’ nature of the non-combatant immunity claim (which possibility I discuss in chapter 5), or else to collapse arguments from innocence into arguments from self defence, which will be examined in the next two chapters.⁴⁵

Treating soldiers as a class

Consider the argument that it is all right to target soldiers, despite the fact that many are innocent, because *most* of them are guilty. There are two slightly different ways of interpreting this argument. The first is to say that because the majority of soldiers in a given army really are morally responsible for their position, and therefore non-innocent, targeting the enemy soldiers *as a class* is a legitimate activity. This is to say that if there are enough legitimate targets within a class, the whole class becomes a legitimate target, and thus even if an attack kills more innocent than non-innocent individuals, that attack is still justified. The second interpretation agrees that the only legitimate targets are the non-innocent soldiers, and that it is as unjustified to kill soldiers who are innocent as it is to kill civilians. However, there are times we think it is acceptable to kill civilians, namely when the doctrine of double effect allows it. The doctrine can also be used to justify attacks on the enemy army, despite the fact that many innocent soldiers will be killed.

I will refer to the first interpretation as the ‘whole-class’ target justification, and the second as the ‘within-class’ target justification.

Whole-class target justification

Whole-class target justification, relying as it does on the idea that having a majority of legitimate targets within a class of persons legitimises treating the whole class as

⁴⁵ If individuals can be rendered immune from attack on grounds of innocence, then in the case of the third class of innocent soldiers discussed, those justly fighting a just war, no attack on them can be morally permitted at all. This is because these soldiers are *all* innocent in the same way as civilians are. My comments in the following discussion therefore apply only to cases in which neither side is justly fighting a just war, but in which members of (at least one of) the other two classes of innocent soldiers are present. As I have discussed in section 3 of chapter 1, the practical implications of this are negligible as both sides always claim to be fighting a just war and will therefore claim that their enemy’s army is not entirely comprised of innocent soldiers.

legitimate targets, is an unviable proposition, and this can be shown with a simple example. Imagine that there are 40 innocent people in a football stadium, being held hostage by 60 terrorists. Would it be acceptable to treat all 100 people as targets and release lethal gas into the stadium, simply because a majority of them were not innocents? Clearly not, but *why* not? One obvious reason is that there are ways to remove the threat of the terrorists more accurately than poison gas, ways that would result in fewer civilian casualties. But note that this is not a reply someone using a whole-class target justification can give. To them, *all* the people in the stadium are now legitimate targets. There is no reason to try to avoid killing any particular subsection of the group. Nor can they say 'it is justifiable to sacrifice 40 innocent lives for the sake of ending 60 guilty ones' – that would be to say that the 40 innocent people were not legitimate targets in themselves. This analogy demonstrates the fundamental problem with a 'majority rules' approach to justifying the attack of certain targets. Being in the middle of a lot of legitimate targets is not enough to make an innocent man a legitimate target in himself.⁴⁶

A response to this objection might be that soldiers are a special sort of class – they are less different from one another than terrorists are from civilians. Soldiers are all engaged in approximately the same activities, they all wear the same uniform and look the same, and they act in support of one another. To target them together is thus justified in a way that targeting civilians and terrorists together is not. I will weigh the merits of these 3 ideas in reverse order. First, consider the idea that it is permissible to target innocent soldiers because they are acting in support of the non-innocent ones – in other words, to stretch the 'non-innocence' justification for killing to apply not just to the non-innocent themselves, but to all those in support of the non-innocent. This is to unacceptably (at least to the immunity theorist) widen the category of those who are permissible targets within war, allowing us to bomb those civilians providing support to the military (and perhaps also those supplying support to *those* civilians, and so on).⁴⁷ Second, the question of appearance: the fact that soldiers fighting

⁴⁶ This is not to say that an innocent man's life must be spared at all costs – there may be cases in which killing him is justified. But that is not the same as alleging that he has become non-innocent and therefore a legitimate target. Rather, in such cases, his death would be seen as regrettable, though necessary. If he was a legitimate target, no such regret would be appropriate.

⁴⁷ My opponent might argue that 'innocent' soldiers supporting the non-innocent within the military are in a different category to civilians offering the same support, because *all* soldiers (innocent or not) pose an immediate threat to the enemy in a way that civilians do not. This argument cannot affect the failure

together in an army all wear the same uniform is straightforwardly and exclusively psychological in its appeal. It is true that the homogeneous appearance of soldiers may make it psychologically easier to target them as if they were all the same, but that is not a moral justification; wearing the uniform of a guilty man does not make John a guilty man. Third and last, we have the claim that all soldiers are engaged in roughly the same activities, unlike the civilians and terrorists, and can therefore all be targeted together. Without further argument, this is to ignore or take to be irrelevant the moral status of the ‘innocent’ soldiers in the class. If we were to agree that soldiers are legitimate targets simply by virtue of engaging in soldiering, then there would be no further argument about the combatant/non-combatant distinction.⁴⁸ This reply, therefore, does not answer the objections to whole-class target justification; it rather denies that objections can be presented, thus ignoring the important moral facts which motivated the ‘innocent soldiers’ problem in the first place.⁴⁹

The reality of war means that it is not in fact possible to target the enemy in such a way that we could be sure that most of the soldiers killed were non-innocents. But someone endorsing a whole-class target justification would be committed to saying that even if it *were* possible, there would be no moral reason to prefer this option. It seems clear, however, that if we take arguments about innocence at all seriously, then we ought to prefer this option. For these reasons the whole-class position is indefensible.

Within-class target justification

Someone holding a within-class target justification position would argue that although it is only morally legitimate to target an army’s non-innocent soldiers,

of a justification *from innocence* for the non-combatant immunity distinction, because it is unrelated to issues of moral innocence. It is closer to an appeal to self defence, which I consider in chapters 3 and 4. If, on the other hand, my opponent argues that ‘innocent’ soldiers, and not civilians, are legitimate targets because of the support they offer the non-innocent soldiers, and gives a reason relating to the activities in which all soldiers are engaged, he begs the question.

⁴⁸ An argument might aim to show that all soldiers are, in some sense, legitimate targets *by definition* (whether this is justified by an appeal to innocence or to self defence), and so it does not make sense to differentiate between them. But if we are not permitted to examine the relevance of the moral concept offered as justification to facts about individual soldiers, then the argument begs the important questions about the distinction under discussion.

⁴⁹ The injustice of extending non-innocence to the innocent soldiers in the class in this way seems compounded when we consider again the reasons we might think them to be innocent, namely because they were in some way coerced into their actions. Like the hostages in the football stadium, many soldiers are part of an army as a result of random factors (e.g. when and where they were born), not because of any morally relevant facts about them or because they consented to being there.

attacking the enemy's military in general is acceptable provided the conditions of the doctrine of double effect are met. Let us revisit the terms of this doctrine, and consider them one by one in the present case:

1. The attack is acceptable in itself as an act of war

Targeting non-innocent soldiers in the enemy's army is certainly as good a candidate as anything for an acceptable act of war. So this first condition is met.

2. The direct effect of the attack is acceptable, e.g. rendering an enemy arms factory inoperable

Killing non-innocent enemy soldiers, as the direct effect of the attack, is acceptable. So the second condition is met.

3. The agent intends only this effect, and only this effect is a means to his end

A military leader could certainly *intend* to kill only non-innocent soldiers in the enemy's army; it is less certain whether this could sensibly be said to be the only means to his end. It is at least arguable that if his end is (something like) removing the threat of the enemy army, then killing a great deal of its soldiers, not just the non-innocent ones, is a causally necessary means to his end. So this third condition could only be met in the unusual case that killing only the enemy's non-innocent soldiers was sufficient to attain one's end.

4. The good of this effect is proportional to the evil incurred, i.e. the good is sufficient to compensate for the deaths of civilians

This condition stipulates that the 'good' of killing non-innocent enemy soldiers must outweigh the evil of killing innocent enemy soldiers. I think it is here that the position becomes pragmatically impossible to justify. The numbers game one would have to play to ensure that more non-innocent than innocent soldiers were killed in an attack strikes me as absurdly complicated. If an army is drafted or its members coerced into their positions, how can the enemy possibly expect to know what percentage of troops

are really in support of the military action? Without such knowledge, it will never be certain that more non-innocent than innocent soldiers are killed in any attack on the enemy army, and so it will never be certain that the proportionality requirement of the doctrine of double effect is met. Note that this is making the generous assumption that the killing of innocent soldiers is 'proportional' if *more than half* of those killed are non-innocent. But certainly when we compare civilian with military casualties in an attack, we require a much greater proportion than this of military casualties for an attack to seem morally justified. It is difficult to imagine how one might identify a morally sound measure of proportionality as it relates to a body count, but whatever it is, it is likely to be tougher for the within-class position than my assumption of a bare majority of non-innocent deaths.

Pragmatism

I conclude from the above discussion that neither whole-class nor within-class target justification can be defended as a justification for killing innocent soldiers. We may conclude that it is *not* the case that it is all right to target soldiers as a class, despite the fact that many are innocent, simply because most of them are guilty. The possibility remains that targeting soldiers as a class, despite the moral disparity between members of that class, can be justified on pragmatic or utilitarian grounds. After all, in order to win a war, one *must* be allowed attack the enemy's soldiers as a class.

There are a number of different ways of interpreting what exactly is being claimed here. It might be simply a true claim about the nature of war: It is true that if one wants to successfully fight a war, one must be allowed to attack the enemy army. But this is not enough to show that such attacks are morally justified. Indeed, if attacking the enemy army cannot be given separate moral justification, this might lead us to conclude that war is necessarily immoral. Another way of interpreting the claim is as some sort of means-end justification: If some wars are truly worth winning, then the acts of war *necessary* to victory must also be justified. It could be persuasively argued that attacking the enemy army is one such act; if this is forbidden, then the good of winning a just war will always be lost. A moral theory which says it is morally right, or even morally *required*, to fight a war, but morally forbidden to do what is necessary to achieve victory, seems to endorse exposing one's own troops to pointless

violence. I will not discuss this further here; the use of utilitarian and means-end justifications for acts of war is discussed comprehensively in chapters five and six.

The third interpretation of the pragmatic claim is this: It is both necessary and justified to attack the enemy army because of certain *morally relevant* facts about what armies are and what they do. In other words, while it is a pragmatic statement that attacking the enemy's soldiers is necessary to win a war, this is true because of facts about those soldiers which themselves morally justify attack. What might these facts be, given that they cannot appeal to the non-innocence of the soldiers? The obvious answer seems to be that soldiers, unlike civilians, pose a threat. Because all soldiers pose a threat to us, whether or not they are responsible for their actions, we can target them all. This is no longer an argument about the non-innocence of soldiers. Rather, it says that even if some soldiers are innocent, they are still candidates for attack for a different reason; and this leads us to the second major justification for the combatant/non-combatant immunity distinction, that of self defence. The next two chapters examine this justification.

A final problem

A final potential problem for arguments from innocence is this: If warring soldiers really have lost the protection usually afforded to their lives and are thus liable to be killed, then why is it that the gratuitous killing of soldiers, killing that goes far beyond military necessity, is usually considered to be wrong? It seems that we still think that a soldier's life is deserving of some degree of protection, and how to account for this is a puzzle for the immunity theorist who is committed to saying that the soldier's non-innocence renders him a legitimate target. After all, his non-innocence is not affected by how many of his comrades have been killed in a given attack. It would seem that the total number of non-innocent soldiers killed in an act of war should be irrelevant to the justice of the act, just as the total numbers of guilty criminals executed under a death-penalty law does not affect the justice of that penalty. That it does not seem to be so stands in need of some explanation.

CHAPTER THREE

Justifying self defence

My purpose in this essay is to ask of non-combatant immunity, ‘What justification for killing will permit the killing of soldiers but prohibit the killing of civilians?’ In the preceding chapter I considered the idea that innocence and guilt provide grounds for killing and preserving life, and argued that even if we accept this justification, we cannot use it to defend non-combatant immunity. In chapter 4 I examine self defence as a justification for the killing of soldiers in war. Again I conclude that, whatever its success in the personal case, self defence cannot be used to draw a distinction between civilians and soldiers. In this chapter, though, I wish to examine the personal case of self defence, and ask a similar question to the one above: ‘What justification for killing will permit the killing of aggressors by their victims, but prohibit the killing of victims by their aggressors?’ In other words, how do we ground the individual permission to kill in self defence?

Despite its general acceptance as part of common-sense morality and long standing in the legal tradition as a justification for homicide, the rightness of killing in self defence has come under close scrutiny in the last decade or so. David Rodin and Suzanne Uniacke have both devoted books to the problem of its justification⁵⁰, and there are many recent articles both defending and attacking the just use of defensive force in a variety of cases.⁵¹ In this chapter I argue that there are good reasons for wishing to retain self defence as a legal defence against homicide, just as there are good reasons for wishing to retain non-combatant immunity as a justification for killing soldiers in war. However, just as with non-combatant immunity, I argue that

⁵⁰ Rodin, D., *War and Self-Defence* (Oxford: Oxford University Press, 2002), and Uniacke, S., *Permissible Killing: The Self defence Justification of Homicide* (Cambridge: Cambridge University Press, 1994).

⁵¹ Articles I have particularly drawn on here include Otsuka, M., ‘Killing the Innocent in Self defence’ (*Philosophy and Public Affairs*, vol. 23, no. 1, 1994, pp. 74-94), and Wasserman, D., ‘Justifying Self-Defense’ (*Philosophy and Public Affairs*, vol. 16, no. 4, 1987, pp. 356-378), although there are many more. The bibliography of Rodin’s ‘War and Self Defence’ provides a comprehensive list.

the classic doctrine of self defence cannot be justified by appeal to the intrinsic properties or individual rights of the agents concerned across the wide range of cases it purports to cover. Rather, a utilitarian defence of the doctrine is more plausible. In arguing for this conclusion I make appeal to many of the particular arguments given by David Rodin concerning non-paradigm cases of self defence and the difficulty in grounding a permission to kill innocent aggressors. I particularly consider Rodin's claim that our right to self defence contains no permission to kill innocent aggressors, and suggest that a utilitarian account of self defence avoids this undesirable conclusion.

Section 1: Necessary and sufficient conditions for self defence

The paradigm military combatant is possessed of both of two properties, either of which might be appealed to as justification for targeting him. He is non-innocent in some important moral sense, and he also poses a threat to his enemy.⁵² The difficulty for the immunity theorist is in explaining why we may always target combatants even when one or both of these properties is absent, and why we may never target non-combatants even when one or both is present. A similar challenge confronts the theorist who seeks to justify our permission to kill in self defence by appeal to intrinsic differences between aggressors and victims. Here, too, the paradigm aggressor in the case of personal case of self defence is non-innocent in some relevant sense, and poses a threat to the victim. However, as I discuss below, there exist many non-paradigm cases in which one of these properties is absent, but in which we traditionally grant defensive permissions to the agent under attack nevertheless. A close examination of these non-paradigm cases highlights the limitations of any deontological account of self defence which appeals to the moral differences between victims and aggressors to justify killing.

Necessity and proportionality

The usual constraints placed on when we may legitimately be said to act in self defence are necessity and proportionality. 'Necessity' refers to the idea that for killing in self defence to be justified, lethal force must be the only way to remove the threat

⁵² I discuss this cross-cut definition at greater length in section 3 of chapter 4.

being faced, e.g. I am not justified in killing a would-be axe-murderer if I can readily drive to safety.⁵³ ‘Proportionality’ refers to the idea that there must be a balance between the harm I inflict in defending myself, and the threat posed to me. Thus, I may not kill you to prevent you stepping on my toe (even if this met the requirement of necessity). One result of these dual constraints is that killings properly described as defensive are automatically limited, both legally and in the philosophical literature, to those cases in which the aggressor poses a serious threat. This is a necessary condition of self defence; the physically healthy adult who kills a helpless little old lady does not do so in self defence, even if the little old lady makes an earnest attempt at inflicting violent harm.⁵⁴ Even in the non-paradigm cases of self defence I discuss below, the ‘victim’ is always being threatened by the ‘aggressor’; that is, the victim’s act of killing always meets the requirements of necessity and proportionality.

Therefore of our two moral properties which might justify killing, moral guilt and the posing of a threat, the latter is constant across all cases of justified self defence. However, as I discuss below, the former is not. Although the aggressor always poses a threat to his victim, there are many cases in which it seems he is in no way guilty for his act of aggression. If we wish to hold that self defence is justified even in such cases, then we must accept that deontological notions of guilt and innocence cannot tell the whole story about what justifies defensive killing. If, on the other hand, we follow Rodin in concluding that self defence is not justified in cases in which the aggressor is not guilty for his attack, we then face epistemological and procedural problems in turning the doctrine into workable law. I argue in section 4 that these problems suggest that our traditional range of defensive permissions ought to be protected by law regardless of their dubious deontological justification. The arguments I offer are formally similar to those I give in chapter 5, in which I discuss the value of non-combatant immunity as a legal convention despite its inability to be explained by deontological moral thinking.

⁵³ There are interesting arguments as to what degree of risk a victim of an aggressive act must take upon themselves to avoid killing in self defence. For example, I am presumably allowed to kill in self defence when there is only a 10% chance that I could safely get to my car and drive away instead; but if there is an 90% chance of safely escaping, I might be obliged to accept the risk and try to make an escape. Space unfortunately constrains me from considering this problem further here.

⁵⁴ As this example indicates, the rules of self defence recognise that the moral guilt of the assailant is not a *sufficient* condition to justify defensive killing. The question I consider below is whether it might, like the presence of threat, be a necessary condition.

Non-paradigm aggressors

There are perhaps three broad classes of non-paradigm ‘aggressors’ which arise in discussions of personal self defence.⁵⁵ The first is the ‘provoked aggressor’, who is (arguably) at least partially excused from his attack. In this case, not only will the aggressor be less guilty than the unprovoked ‘purely aggressive’ aggressor, but his victim will not be as innocent of the attack as the paradigm victim. I will not consider this case further as I am concerned with self defence specifically as a justification for killing, and I think it unlikely that any degree of provocation is sufficient to justify a lethal attack on somebody.⁵⁶

The second class of non-paradigm aggressor is the ‘innocent aggressor’ who is acting under non-culpable mistake, duress or coercion to the extent that his personal responsibility for his attack is nil. For example, the man who is hypnotised by an evil Svengali and ordered to assassinate you is excused from all moral guilt by the non-voluntary nature of his actions. Even actions which are performed voluntarily may be excused, for example the actor who reasonably believes he is firing a stage-gun loaded with blanks, but which his victim somehow knows is loaded with live ammunition; or the policeman who shoots you because he mistakes you for your armed and dangerous identical twin.⁵⁷ These aggressors are not culpable for the threats they pose, even though they are a direct result of voluntary and intentional actions. In each of these cases, the aggressor is deemed not responsible for the action he performs and/or its consequences, and as he is not responsible, he is not guilty.

The third class of non-paradigm aggressor is the ‘innocent threat’, who poses a threat to the victim not as an intentional actor, but as a mere physical object. One of the

⁵⁵ I place the word ‘aggressors’ in quotation marks because, like ‘combatant’, it is a technical term which picks out a diverse class of persons, many of whom do not fit the intuitive characterisation of an aggressor. In the following discussion, ‘aggressors’ are all those who pose threats, and ‘victims’ are all those who are threatened.

⁵⁶ It should be noted, however, that in cases where someone acts to defend property or to repel a less severe physical assault, considerations of provocation have more weight. We certainly have mixed sympathy for the ‘victim’ who returns the punches of an aggressor he has goaded into attacking him, even if we consider that aggressor unjustified. It seems that some kinds of provocation mitigate a permission to defend oneself.

⁵⁷ It is important to remember that these killings will only be permitted if they are necessary and proportionate, so the policeman acts non-culpably only if he would have been permitted to attack your twin in like circumstances.

most cited (and least plausible) examples in the literature comes from Robert Nozick and concerns a villain who grabs a fat passer-by and pushes him down a narrow well onto his (the villain's) enemy, who is trapped at the bottom.⁵⁸ The fat man is an aggressor only insofar as his physical bulk poses a lethal threat to the man at the bottom of the well, and it is generally agreed that he (the fat man) is entirely innocent. Other examples include an epileptic person having a fit who pushes other pedestrians into the path of oncoming traffic, or a hiker who stumbles on a high narrow mountain path, unbalancing her companions. These people pose threats to others, but they do not truly aggress. For the purposes of this discussion, however, I will use the term 'aggressor' to include these cases.

Provided we adhere to the limits of proportionality and necessity, are we permitted to kill to defend ourselves against these two latter types of 'aggressors'? It is this question which I consider below. Any answer to this question will ultimately depend on acceptance of a specific moral theory, and I cannot hope to exhaustively explain here why the theoretical basis for one answer ought to be preferred over that of another. However, I hope to show four things: First, our traditional permission to kill innocent aggressors cannot be justified by Kantian or rights-based arguments. Second, a wide range of defensive permissions can be justified by appeal to utilitarian considerations, including the permission to kill innocent aggressors. Third, when considering self defence as a legal doctrine, there are strong reasons for wishing to retain this wide range of defensive permissions. Fourth, it is an advantage of the utilitarian position that it can reconcile the desirable legal doctrine of self defence with the moral doctrine.

In the following section I consider how a rights theorist and a Kantian might explain the right to kill in self defence, before turning in section 3 to a utilitarian account of the right both as a legal and as a moral doctrine. Following the arguments given by David Rodin in *War and Self-Defense*,⁵⁹ I explain why rights theorists and Kantians are both committed to the conclusion that we have no moral permission to kill innocent aggressors in self defence. That is, when we knowingly and intentionally kill

⁵⁸ Nozick, R., *Anarchy, State and Utopia* (Oxford: Blackwell, 1974), p. 34

⁵⁹ Rodin, D., p. 35ff

such people ‘in self defence’, our act is not one of justified self defence but rather one of culpable homicide. In section 3 I explain why a utilitarian account of self defence does allow us to defend ourselves against innocent aggressors much of the time, although it will not grant an absolute permission to do so. In section 4 I examine self defence in the context of legal convention, and show how the utilitarian can offer a unified account of the legal and moral doctrines of self defence.

Section 2: Explaining deontological defensive permissions

Rights theory

In *War and Self-Defense*, David Rodin gives a thorough and persuasive discussion of how the personal right to self defence might be grounded in rights theory.⁶⁰ He begins by explaining the ‘default’ position with regard to our rights: we all possess a right to life which includes a claim against others that they not kill us. He then explains that we are permitted to kill an aggressor in self defence because the aggressor acts in such a way that he loses his claim on us that we not kill him, by losing his right to life with respect to us:

The aggressor’s right to life includes the claim against others that they not kill him. If the aggressor forfeits this right with respect to the defender then he has a no-claim against the defender that he not kill him.... [The victim’s] right to kill and the absence of the aggressor’s right to life are just the same moral fact described in different ways.⁶¹

The explanation for this ‘moral fact’ can be found, says Rodin, in ‘the fault of the aggressor for the aggressive attack.’⁶² In other words, the aggressor loses his right to life with respect to a certain victim if and only if he (the aggressor) is at fault for his attack on that victim. Rodin gives a step-by-step explanation of the interaction between rights:

⁶⁰ Rodin, pp. 70-102

⁶¹ Rodin, p. 75-6

⁶² Rodin, p. 77

I have the right to life. Therefore, if an aggressor makes an attack on my life, in the absence of special justifying circumstances, he wrongs me. *Because I am innocent and he is at fault* for the aggression, his claim against me that I not use necessary and proportionate lethal force against him becomes forfeited (or fails to be entailed by his right to life). Therefore I have a right (liberty) to kill him. Therefore when I attack him to defend myself, I do not violate his right to life, and hence I do not wrong him. Because I do not wrong him, I do not forfeit (or fail to possess) my right to life.⁶³

To put it briefly, Rodin is saying that the asymmetry in the rights of the victim and the aggressor is a result of their asymmetry in respect to moral guilt and innocence. This explanation is intuitively extremely appealing, both in its simplicity and because, as I have discussed in chapter 2, we do usually think that guilt and innocence ought to affect how agents are treated. However, as Rodin acknowledges, this explanation does not provide a justification for self defence when we are attacked by innocent aggressors or innocent threats. Such people, in common with the general population, still hold a right to life which includes a claim against us that we not kill them.

Kantian moral thinking

As well as their possession of a right to life and corresponding claim that we not kill them, Rodin invokes a Kantian notion to explain why we may not kill innocent aggressors: the idea of respect for persons. In attacking someone who culpably and voluntarily attacks us, Rodin says, we are manifesting an entirely appropriate and respectful attitude to that person as a moral agent. The threat he poses derives from him ‘as a moral subject, not just as a physical entity’, and in fighting back we treat him as a subject and not merely as an object.⁶⁴ However, when we fight against innocent aggressors, it is no longer the case that our actions are respectful of that agent’s personhood:

The only reasons we could have for taking [innocent aggressors’] lives is that they happen to be situated in such a way that their death is necessary

⁶³ Rodin, p. 79, italics added.

⁶⁴ Rodin, p. 88

for our survival. But this is not a relevant fact about them as subjects and harming them on this basis could manifest no proper attitude toward them as persons.⁶⁵

A similar appeal might be made to the Kantian prohibition on treating others as mere means to one's end. Sacrificing the life of an innocent person in order to protect one's own is a grievous example of treating them as a mere means, and killing an innocent aggressor amounts to this kind of sacrifice. So, like the rights theorist, the Kantian seems committed to the view that it is wrong to kill the innocent in self defence.⁶⁶

In summary, non-paradigm cases of self defence involving innocent aggressors and innocent threats suggest that there is no feature intrinsic to all 'aggressors' which justifies killing them. Sometimes the fact that someone is engaged as an aggressor in a self defence situation is in no way a reflection on him as a moral being, and he still retains the 'default' protections accorded to any innocent person. (What the precise nature of these protections is taken to be will vary from theory to theory, but every theory will presumably wish to treat like cases alike.) There remain two avenues of justification for the traditional broad range of defensive permissions. First, killing in self defence might be justified by some deontological consideration which appeals exclusively to the victim's innocence. That is, we might each have an absolute right to take whatever proportionate means necessary to defend our own lives, even when this requires the killing of innocents, so long as we are ourselves are not responsible for the situation which forces us to choose between lives. I consider this below. The second possibility is that the doctrine can be justified by extrinsic features of the parties involved; that is, by appeal to consequentialist considerations. I discuss this further in sections 3 and 4.

Arguments from the victim's innocence and the bystander problem

⁶⁵ Rodin, p. 89

⁶⁶ Barbara Herman concurs. She gives a Kantian account of self defence in *The Practice of Moral Judgment* (Cambridge, Ma: Harvard University Press, 1993), and concludes: 'The argument that justifies violence in self-defense does not justify violence against innocent threats... Self-defense is permitted as a way of resisting a willed attack on my agency. Since the innocent threat is identified as having no agency-discounting maxim – no maxim of aggression – I may not act against him in self-defense.' (p. 130)

While there might sometimes be a utilitarian or even a deontological justification for sacrificing an innocent in order to save oneself, it is hard to imagine how to ground in deontological moral theory the claim that each innocent person has the *inalienable* right to kill, when necessary and proportionate, in order to stay alive. The problem, as Rodin and others⁶⁷ have pointed out, is that to accept this robust permission in cases of self defence obliges us to accept it in a wide range of other ‘either him or me’ cases in which the mainstream deontological moral theorist will generally consider it wrong to kill. Consider the following scenario: I am in a crowd and a villain, mistaking me for his arch-enemy, throws a javelin at me. Lightning-fast, I grab a random bystander and hold him in front of me, protecting myself from the javelin’s impact but sacrificing the bystander’s life thereby. Now, insofar as my act is both necessary and proportionate, and in terms of the intrinsic properties of the parties concerned, this case seems morally parallel to the case in which I kill an innocent aggressor. In each case, I have to make an immediate choice between saving someone else’s life and saving my own. In each case, I am in no way culpable for the situation which makes this necessary, and nor is the person whose life I sacrifice. Granting that the two cases really are analogous, then, any ‘either him or me’ right to kill innocent aggressors will presumably also be a right to kill innocent bystanders. It may also be a right to kill someone to take their liver for an emergency transplant, or to kill an innocent person when a villain threatens to kill *you* if you do not; and many more killings of this ‘either him or me’ sort.⁶⁸

Now, I have not shown that it is impossible to ground a deontological moral principle that ‘anything necessary and proportionate is permitted when it’s him or me’. I have considered only one of the many arguments a deontologist might give for the right to kill in a broad range of self defence situations. However, it seems that whatever

⁶⁷ As well as Rodin (p. 81), Michael Otsuka discusses the below-mentioned analogy between killing innocent aggressors and killing bystanders at some length in ‘Killing the Innocent in Self-Defence’ (*Philosophy and Public Affairs*, vol. 23, no. 1, 1994, pp. 74-94).

⁶⁸ The deontologist may wish to object that such killings are not proportionate, perhaps for the same reasons as I offer in my later discussion of the utilitarian permission to kill. However, the deontologist (unlike the utilitarian) seems committed to the idea that human lives are of equal intrinsic value, and thus exchanging one for another is always proportionate. This is certainly Rodin’s view: in his thorough rights-based discussion of the proportionality requirement of self defence, he states that “‘Life for life’ is a very clear and uncontroversial judgement, and it provides the basis for our understanding of proportionality.” (Rodin, p. 43. He also allows that one may kill to defend those things on which one’s life depends, that is, I may kill a thief to prevent him stealing my food supply, p. 44.) Thus the deontologist who objects that my killing an innocent bystander to save myself is not *proportionate* killing must explain why the bystander’s life is more valuable than my own.

argument she offers, if she is consistent, the deontologist will be forced to the conclusions I indicate above. She may, of course, simply bite the bullet when it comes to permitting killing in these other cases, but the bullet is a hard one to bite for any theorist purportedly committed to the preservation of innocent lives. It strikes me that the more attractive option for the deontologist is to accept that, tradition notwithstanding, we have no permission to defend ourselves against the attacks of innocent aggressors. This is the conclusion which Rodin adopts and defends.

There remains the possibility that a permission to defend against innocent aggressors can be justified by appeal to non-deontological notions. I present an argument for this conclusion in the next section, and follow it in section 4 with an argument that it is desirable to retain a broad range of defensive permissions in law, regardless of their lack of deontological justification.

Section 3: The utility of ‘either him or me’ killing

So far I have restricted my discussion to deontological accounts of defensive rights. I have argued, first, that innocent aggressors possess no intrinsic qualities which might justify killing them; and second, that any deontological argument for a right to kill innocent aggressors will also be an argument for a right to kill innocents in many non-self-defence cases, such as the innocent bystander case. Taken together, these arguments suggest that there is no adequate deontological account of our traditional broad range of defensive permissions which does not license still broader permissions to kill.

In this section I explain why a permission to take whatever proportionate means necessary to defend our own lives fares much better on a utilitarian framework. I argue that in the majority of self defence cases, including those involving innocent aggressors and innocent threats, utility is best maximised by granting each innocent person permission to kill to defend her life; but when this is not the case, no such permission exists. Thus utilitarians will often, but not always, permit the killing of innocent aggressors and innocent threats (and sometimes innocent bystanders as well). In section 4 I explain why there are utilitarian reasons for thinking it desirable

that the legal doctrine of self defence permit the killing of innocent aggressors, and suggest that the challenge for the deontologist is to allow for the evident desirability of this legal doctrine of self defence despite its divergence from the moral one.

Foreseeable consequences and self defence

Here again is the principle which seems necessary to justify the broad range of defensive permissions we traditionally and legally hold, but which I have suggested the deontologist might find unpalatable and possibly difficult to justify:

We are permitted to take whatever proportionate means necessary to defend our own lives, even when this requires the killing of innocents, so long as we are ourselves are not responsible for the situation which forces us to choose between lives.

Compare this to the utilitarian standard of permissible defensive killing, which could be formulated like this:

We may kill others to save ourselves just as long as the foreseeable consequences of this action are comparable to or better than those of other possible actions.⁶⁹

I argue that in the vast majority of actual cases in which an act of killing is necessary and proportionate, it will also be the case that the foreseeable consequences of killing one's assailant are preferable or comparable to the foreseeable consequences of dying oneself. In other words, instances of 'justified self defence' as we traditionally understand it are overwhelmingly likely to meet the utilitarian standard of permissible defensive killing. Moreover, given that human beings have a severely limited ability to make the sort of split-second expected-utility calculations that self defence requires, and must in such cases employ utilitarian 'rules of thumb' which permit

⁶⁹ I am making here two assumptions about utilitarian moral theory. First, the permissibility of an action is decided by that action's foreseeable consequences, not by its actual consequences. Second, when an agent chooses between two actions which maximise utility equally, and to a greater degree than any other action, then the agent is right if she performs *either* maximal action (and wrong if she performs any non-maximal action). I take this to be a straightforward consequence of the utilitarian standard of rightness. For corroboration of this interpretation see Smart, J. J. C and B. Williams, 'Utilitarianism: For and Against' (Cambridge: CUP, 1973), p. 45.

defensive killing (see below), I consider it unlikely that any plausible self defence situation will be of a kind in which the innocent victim does wrong to defend herself against a lethal attack made by an innocent aggressor.

I do not mean to claim that utility is never *in actual fact* maximised by an innocent victim allowing herself to be killed by an innocent aggressor. We can imagine a case in which a man who knows the cure for all cancers, but carelessly hasn't yet told anyone or written it down, is induced by hypnosis to attack an innocent victim. This victim, in defending herself, will fail to maximise utility (assuming that this innocent aggressor plans to put his knowledge to good use). However, it does not follow that she has done something morally wrong. What makes her action right or wrong is its foreseeable consequences, and the victim in this case cannot foresee that in killing her assailant, she is destroying valuable medical knowledge. All she knows is that she, an innocent person, is being attacked by a (she supposes) homicidal maniac, and she is reasonable in assuming (if she has time to reflect on it at all) that to preserve her own life in preference to that of her assailant will not fail to maximise utility. Indeed anyone who is the victim of an (as far as they can see) unprovoked attack will very probably receive the impression that their life is more worthy of protection, in terms of expected utility, than that of their attacker.⁷⁰

'Rules of thumb'

As I hinted above, it may often be that an agent defending against an unprovoked lethal attack is not able to make any sort of calculation at all about the expected utility of her actions. In this case, she can increase her chance of maximising utility by following a 'rule of thumb' which generally maximises utility. As I discuss further in section 4, a rule permitting killing in self defence is one of these. J. J. C. Smart argues for the role of 'rules of thumb' in utilitarianism in *Utilitarianism: For and Against*.⁷¹ He illustrates his argument with the example of a split-second decision to rescue a

⁷⁰ I am making certain general assumptions here about the average 'innocent victim'. It is entirely possible that an evil scientist plotting to destroy the world could be the victim of an unprovoked attack, and in such a case the scientist could not reasonably expect to maximise utility by preserving his own life (unless he also intended to mend his ways). Such a person would, if attacked, be wrong to kill to defend himself. By the same reasoning, the victim who is attacked by the man who knows the cure for cancer (discussed above) might, if she was aware of her attacker's secret knowledge, have an obligation to sacrifice herself.

⁷¹ Smart and Williams, pp. 42ff.

drowning man, which is in some ways comparable (in the urgency of the situation and the strong natural impulses governing our decisions) to the case of self defence:

[The utilitarian theorist] knows that a man about to save a drowning person has no time to consider various possibilities, such as that the drowning person is a dangerous criminal who will cause death and destruction, or that he is suffering from a painful and incapacitating disease from which death would be a merciful release, or that various timid people, watching from the bank, will suffer a heart attack if they see anyone else in the water. No, he knows that it is almost always right to save a drowning man, and in he goes.⁷²

Likewise, we know that it is almost always permissible to kill an unprovoked assailant in order to defend one's own life. When we are prevented, by time constraints or some other impairment, from performing any calculation of expected utility, we do right to follow this rule of thumb.⁷³

Hard cases for the utilitarian?

I stated above that anyone who is attacked by an apparent homicidal maniac does no wrong to judge (assuming she has leisure to) that defending herself with lethal force will maximise utility. However, there are many 'either him or me' cases in which the victim has no reason to suppose that her assailant is a homicidal maniac; what of these cases? Some of these do not involve self defence as we traditionally understand it, and yet the standards of permissible defensive killing outlined above are met. Examples are the case in which A must kill B in order to prevent C killing A; the case of the innocent bystander mentioned above; and the case in which someone needs an emergency organ transplant.

One problem I suggested for a deontological principle allowing 'either him or me' killing is that it seemed to allow us to always kill in these situations, despite the

⁷² Smart and Williams, p. 43

⁷³ Smart explains that what is required is for the utilitarian to 'school himself to act . . . habitually and in accordance with stereotyped rules', in other words, to endeavour to conform one's unconscious decision-making to utilitarian standards (p. 43). It is fortunate that so many of our instinctive decisions (to kill in self defence, to rescue drowning men, to keep casual promises, to tell the truth etc.) conform to this standard as it is.

traditional view that such killing is often unacceptable. I believe that the utilitarian standard of permissible killing can avoid this conclusion. This is not because it will infallibly prohibit killing in such cases; it will not. It may often be, as I explained above, that any calculation of utility is impossible, and in these cases the 'rule of thumb' which states that killing innocent people generally fails to maximise utility must come into play. However, where calculation is possible and killing does seem permitted, the justification for permitting the killing will be identical to the justification for generally forbidding it, i.e. the maximisation of expected utility. In this way the utilitarian is consistent in her preservation of what she professes to value, in a way which may be problematic for the deontologist who wishes to protect the lives of innocents. I discuss these 'either him or me' hard cases below.

The utility of sacrificing innocents for innocents

The utilitarian states that a victim may choose to kill rather than to die as long as she reasonably judges that her act of killing has consequences which are better than *or comparable to* that of her act of submission to death. It need not be that the aggressor is a villainous rogue who is likely to cause massive amounts of disutility to the world. It merely needs to hold that the foreseeable consequences of the aggressor's death are no worse than the foreseeable consequences of the victim's death. This seems to permit killing in a far wider range of cases than the 'apparent homicidal maniac' situation I discussed above. For example, if two innocent people were tussling on a dangerously narrow mountain path, each might reasonably foresee that her own death, in terms of expected consequences, would be of equal disvalue to that of her companion. Thus each would be permitted to try to save herself, even if this meant pushing her companion off the path.

In this way, the utilitarian permission to kill in self defence is more extensive than that found in Kantian or rights-theory; but it is more restrictive than the 'anything goes when it's either him or me' principle considered above. For example, to the utilitarian I am probably never permitted to kill somebody to take their organs, even if the alternative is that I will myself die. This is because it is probably impossible to truthfully claim that I reasonably expect to maximise utility by killing another person and taking their organs. Unlike the somewhat artificial and isolated mountain-path example above, the real-life consequences of forced-organ-sacrifice are tremendously

complicated. Organ transplant necessarily involves the co-operation of many people with specialised skills, high post-operative risk, expensive medical care, and so on.⁷⁴ Moreover, even if an instance of secret (so as to avoid the negative consequences of social fear) organ-sacrifice were to maximise utility when considered in isolation, it must be remembered that the required secrecy is probably not possible within a healthcare system which itself succeeds in maximising utility. That is, the conditions necessary for organ-sacrifice to maximise utility themselves fail to maximise utility and thus could never be permissibly implemented.⁷⁵

So much for organ sacrifice and similarly complex situations. More difficult, because potentially less complex, is the case in which innocent man A is told by villain C that he, A, must kill innocent man B, or else C will kill A. In such a case, the utilitarian can say this: First, 'kill or be killed' is not always an exclusive description of the actions available to an agent; in many plausible situations of this sort, utility will be maximised by A's appealing to the police. Second, if A reasonably expects that if he himself does not kill B, someone else will be persuaded to do so, then *ceteris paribus* he is permitted to kill B. To refuse to do so at the expense of his own life is merely to add another name to the list of casualties of C's villainy. Even if A is of the opinion that B is likely to be spared by C if this attempt on him fails, A may choose to kill B provided he has reason to suppose that the consequences of this action are preferable to or comparable with the consequences of A's death. However, this utilitarian

⁷⁴ Of course, the utilitarian is committed to saying that if all these factors were accounted for and it were certain that utility would in fact be maximised in this way, then killing somebody to take their organs would be permissible. However, this bullet is an easy one to bite, because the antecedent conditions necessary for such a moral judgement to be made are so implausible. In other words, in most cases we will not be able to calculate carefully, but will instead rely on our utilitarian 'rule of thumb' which states that it is nearly always wrong to kill someone to take their organs, because this action nearly always fails to maximise utility.

⁷⁵ This consideration will not affect the utilitarian who is living in a world in which healthcare systems are already corrupt and secretive. But it is nevertheless worth noting that in order for organ-sacrifice to be permissible to the utilitarian, we must have already failed to maximise utility in a number of ways. In fact, many acts the utilitarian is vilified for permitting will only be permitted in already-undesirable situations. For example, as I discuss further in section 5, it will sometimes maximise utility to kill an innocent soldier; but it only does so in virtue of others' previous or anticipated non-utility-maximising behaviour. I do not consider it a fault of utilitarianism that it admits or requires usually-impermissible acts against a backdrop of its own misapplication. It has a good explanation for why the act is now permissible: because all preferable alternatives have been eliminated by un-utilitarian behaviour. (Contrast this with a classic Kantian intuitive problem, that we have an obligation to tell the truth even when a homicidal maniac asks if we know where his victim is to be found. The traditional Kantian response is that the badness of the situation is caused by other people's moral wrong-doing, and so they are responsible for the outcome; but many people find this unsatisfactory. It seems more appealing to say that even if others' faulty conduct creates some bad situation, we ought not to use that as an excuse to allow a still worse situation to develop.)

standard will be more difficult to meet when it is taken into account that in killing B, A is sending C a strong message that C's strategy for removing his enemies is a good one, thus encouraging its repetition and generating much future disutility. I am reminded here of the position of the UK government in Iraq in October 2004 with reference to British hostage Ken Bigley.⁷⁶ The demands of Ken Bigley's hostage-takers seemed minor, particularly when balanced against his (substantially) innocent life. However, the British government argued that granting even a minor demand to secure the release of one man would inevitably lead to an increase in the practice of hostage-taking, and thus put the lives of many more people in danger.⁷⁷ The same consideration could affect the utility of A's decision to comply with a demand to kill B on behalf of C.

It is worth stressing once again that these examples, in their exclusive 'choice 1 or choice 2' format, are unlikely to be similar to real life cases of self defence of other life-or-death choices. It is not easy to think of a plausible case of personal self defence in which the aggressor is innocent, and those we can think of tend to have such complicated foreseeable consequences that one might say that for the victim to make any kind of split-second moral judgement is simply impossible.⁷⁸ Having said this, the discussion above suggests that a utilitarian account of permissible killing will often accommodate the broad range of defensive permissions found in our moral tradition. I will now explain why there are good pragmatic and utilitarian reasons to wish to retain this broad range of permissions in law, regardless of their intrinsic justification.

Section 4: Self defence in law

⁷⁶ Information about Ken Bigley's killing are widespread on the internet, but a good summary of the story can be found on the Washington Post website, October 9, 2004. Retrieved from www.washingtonpost.com/wp-dyn/articles/A17349-2004Oct8.html.

⁷⁷ British Foreign Secretary Jack Straw stated: 'We cannot get into a situation, and I believe the family understand this, where we start bargaining with terrorists and kidnappers, because if we were to, we would not make Iraq or anywhere else safer. We would make everywhere much less safe.' (Newsmax story of December 200, 2004, retrieved from www.newsmax.com/archives/articles/2004/9/22/162522.shtml.)

⁷⁸ Perhaps the most obvious cases of the killing of innocent aggressors occur in war, with the apparently defensive killing of innocent soldiers. However, as I discuss in the following chapter, it is unusual for the conditions of proportionality and necessity to be met in these cases, and so they are not properly speaking cases of individual self defence.

In examining how self defence is justified in the personal case, I have been primarily concerned not with the legal doctrine of self defence, but with the moral one. This is because of the role played by self defence in arguments about non-combatant immunity: it is often thought that we are morally permitted to kill the enemy's soldiers, be they innocent or guilty, just because they pose a threat to us. Deontological justification for these acts of killing thus relies on us having certain kinds of deontological defensive permissions. However, if there is no way of grounding a deontological moral permission to kill innocents in self defence, and I have argued above that there is not, then it follows that such arguments cannot sustain a permission to kill innocent soldiers. Thus the immunity theorist must either relinquish her claim to deontological justification, or find some other deontological moral feature common to soldiers which justifies killing them. This is why the justification of self defence as a moral doctrine has implications for the justification of non-combatant immunity as a moral doctrine.

The role played by self defence as a legal doctrine is also important, however, both in considering that doctrine's potential ultimate justification, and to the arguments I give concerning non-combatant immunity. As I argue in chapter 5, despite the lack of deontological justification for the distinction drawn between soldiers and civilians, there are excellent utilitarian reasons for wishing to retain non-combatant immunity as part of *jus in bello*, at least until a more progressive doctrine is enforceable. Likewise, I believe that there are excellent pragmatic reasons for wishing to retain a doctrine of self defence which permits the killing of innocent aggressors, despite the fact that there is no deontological justification for this type of killing. As I argue below, these pragmatic reasons are also utilitarian reasons, and thus the utilitarian can subsume the legal doctrine of self defence under the moral doctrine. The deontologist, however, if she wishes to avoid radically unpragmatic conclusions regarding a legal right to self defence, seems committed to separate justifications for the legal and moral doctrines of self defence. I hold it to be an advantage of utilitarian moral theory that it can offer a unified argument for, and explanation of, each doctrine.

A legal permission to kill innocent aggressors

There are two chief reasons we ought to wish to retain the legal permission to defend ourselves against innocent aggressors. First, if my potential aggressor knows that I,

his proposed victim, entertain a robust permission to kill him in self defence, the deterrent effect of this possible outcome is maximised. Second, human epistemic limitations mean that it will generally be impossible to judge whether a victim knew his aggressor was innocent. The usual approach in law is to apply a 'reasonable man' standard: the victim is assumed to have been able to make whatever judgement a 'reasonable man' would in his place and is judged accordingly. However, the 'reasonable man' standard is undermined in self defence cases by our conviction that in urgent life-or-death decisions of this kind, the 'reasonable man' will probably fight for his life regardless of his beliefs about the culpability of who is attacking him.

The nature of law requires that legal doctrines make broad judgements which are held to apply across a range of cases, and which create a desirable situation when so applied. When it comes to in self defence, the relevant broad judgement is that an individual may, when necessary and proportionate, kill in self defence. This is a statement of the utilitarian 'rule of thumb' about killing in self defence. Any limitation placed on this judgement concerning the innocence of the aggressor invites a welter of epistemological problems with which the legal system (and indeed any social system) is ill-equipped to deal.

Because its moral permissions approximate existing legal permissions, the utilitarian account of self defence can give us an explanation of why the legal doctrine of self defence ought to be as it is: taking into account the limitations of the legal system and of human epistemic powers, it is maximally successful at generating utility. Deontological moral theory, on the other hand, cannot provide a correspondence between the legal and moral doctrines of self defence without making important changes to the legal doctrine forbidding the defensive killing of innocents. For these changes to be enforceable, we would require a solution to what I suggest may be insoluble problems arising from human epistemic limitations. The deontologist seems committed to the uncomfortable view that what is morally desirable is pragmatically undesirable.⁷⁹

⁷⁹ This is not an unusual situation for the law to be in. There are many moral doctrines which would make terrible law, and we are not troubled by this fact. A law against rudeness, or against the minor telling of lies, would be pragmatically unfeasible and undesirable. However, the gap between the legal and moral ought, we usually think, to close when it comes to questions of life or death. If some variety of killing is wrong, it seems reasonable to wish it to be illegal. Whether from motives of retribution or

Section 5: Conclusions

I hope in this chapter to have established three things. First, despite the received wisdom about the robust ‘right’ to self defence, there is no deontological justification for the claim that we may kill aggressors who are not at fault for the threat they pose. Second, a utilitarian account of permissible killing will generally grant both these defensive permissions, but will not advocate an absolute ‘right’ of the sort deontologists defend. Rather, the permissibility of killing an innocent aggressor depends on the foreseeable consequences of that act of killing. I have suggested that a ‘rule of thumb’ that we may kill in self defence will, in general, generate utility, and so we are right to follow this rule when it is impossible to make more specific utility calculations. Third, there are good pragmatic reasons to retain a broad range of defensive permission in law, and it is a benefit of utilitarian moral theory that the desirability of the legal doctrine can be explained by the moral doctrine.

I now return to my broader project of examining justification of the doctrine of non-combatant immunity.

Self defence and immunity theory: comparing doctrines

Self defence theorists have traditionally referred to ‘victims’ and ‘aggressors’ as if they were classes of persons, just as immunity theorists refer to ‘combatants’ and ‘non-combatants’. In this chapter I hope to have shown that in the former case, these labels are misleading and morally irrelevant. The range of deontologically-morally relevant properties possessed by ‘victims’ and ‘aggressors’ means that any categorical distinction made between them will ride roughshod over important differences between individual cases.

The overall project of this essay is explain why the labels of ‘combatant’ and ‘non-combatant’ are likewise misleading and morally irrelevant, as the classes identified are no more uniform (and in real life, probably much less uniform) than those of

deterrence, we usually think the law ought to be concerned with punishing wrongdoing of this magnitude.

‘victim’ and ‘aggressor’. The doctrine of non-combatant immunity and the doctrine of self defence thus encounter similar problems in their search for deontological justification. They also share another similarity: as I discuss in chapter 5, although it is incapable of deontological justification, there are excellent reasons to retain non-combatant immunity as a legal doctrine.

A return to immunity theory

In the next chapter I return to my discussion of civilian immunity, and examine self defence as it is used to discriminate between combatants and non-combatants in times of war. I use my conclusions from this chapter to argue that, in any case where the deontological justificatory conditions for self defence obtain, there is no reason to suppose that the victims and aggressors are more likely to be soldiers than civilians. I also argue that, due to the requirements of necessity and proportionality, it is very rarely going to be the case that a soldier’s act of killing in war is analogous to the personal case of self defence. More often these considerations will transform self defence into some kind of national-defence – and although some such argument may justify the waging of war, it cannot distinguish between soldiers and civilians therein.

I then examine the idea of formulating a war-specific account of defence, in which some particular degree of threat posed by the enemy is deemed sufficient to justify attack. I explain why this approach, although in many respects a promising method for characterising military necessity, again cannot distinguish between civilians and soldiers. I conclude that arguments from self defence cannot be used to justify the principle of non-combatant immunity.

CHAPTER FOUR

Self defence and war

As I discussed in chapter 2, an argument from innocence for the distinction between combatants and non-combatants is attractive because ‘innocence’ is a notion of accepted moral importance. The appeal of the second major variety of argument used by immunity theorists, arguments from self defence, is that we already recognise, both in our everyday judgements and in judgements of law, that defending oneself is a justified activity. In chapter 3 I explored the justifications for this view, arguing that we ought to endorse a conception of self defence which allows the killing of both culpable and innocent aggressors when consequentialist considerations permit it. In my conclusion I also claimed that due to the requirements of necessity and proportionality, it is very rarely going to be the case that a soldier’s act of killing in war is analogous to an act of killing in the personal case of self defence. I spend the first part of this chapter explaining why I think this is. In the second half I explain why I think that, no matter what type of self-defence justification is appealed to and no matter what degree of threat is deemed proportionate to the use of lethal force, self defence cannot separate the killing of combatants from the killing of non-combatants in the way that the immunity theorist desires. I also explore and reject the possibility of reinterpreting our definitions of ‘innocence’ and ‘self defence’ in order to avoid this conclusion.

Throughout this chapter it must be remembered that we are discussing self defence as applied to specific acts of war, not as the justification for going to war. Defence as an *ad bellum* justification is (in principle, if not in practice) relatively uncontroversial: a genuinely defensive war, in which an aggressor is repelled in order to protect something obviously worth protecting, is seen by many as the best candidate for a just cause that there is. Many people consider World War II to have been a war fought with this kind of just cause (though whether it was justly fought is the subject of intense debate). But an appeal to defensive permissions *ad bellum* does not have any

bearing on self defence as it is used to defend a distinction between civilians and soldiers *in bello*. We must separate the ideas of self (or national) defence providing a just cause for war, and self (or national) defence providing a reason to target only combatants within that war.

Section 1: The conditions for self defence

As my discussion in chapter 3 illustrates, there is disagreement amongst philosophers about some of the conditions which must obtain for killing in self defence to be justified. For example, Rodin argues that the aggressor must be culpable; I argue he need not. But there is general agreement about at least two conditions of justified self defence: the act of killing must be *necessary* in order to deflect the threat, and it must be *proportionate* to the harm posed. The following discussion concentrates on these conditions as they are (as far as I know) uncontroversial, and I believe they are sufficient to illustrate why what justifies self defence in the personal case will rarely justify it in war.

Proportionality

The question of how to properly view the requirement of proportionality in war is somewhat more difficult than in the personal case of self defence. In the personal case, the harm inflicted by the defensive act (i.e. the killing of the aggressor) must be in proportion to the threat posed by that aggressor. It is generally accepted that to kill in defence of one's own or another's life is a proportionate use of force, as is killing in order to prevent oneself or another being permanently maimed. There is still controversy over whether one may kill to prevent rape, or to protect one's property rights.⁸⁰

When it comes to war, however, the question becomes even more vexed, because it is not an easy matter to arrive at a proper conception of what is being defended by a soldier who kills in war. For the analogy to the personal case of self defence to hold straightforwardly, soldier A's act of killing soldier B should be seen as being simply in defence of A's life. However, when soldiers attack one another in war they

⁸⁰ Rodin provides a good discussion of these questions, p. 43ff, as does Uniacke, p. 132ff.

generally do so not just in defence of themselves, but rather in defence of some broader objective.⁸¹ A kills B in order to protect the military campaign in which he is currently employed, and ultimately the cause for which he is fighting: his country's territorial integrity, his political freedom, his cultural autonomy etc. If we understand an act of self defence in war as having this end, then the act will be proportionate only if the good being protected is worthy of this level of protection – and indeed, only if we allow that the analogy of self defence stretches to an individual killing other individuals to protect these types of collective goods. This would mean that individuals *en masse* could kill other individuals *en masse* in a sort of orchestrated act of national-defence, and the justification would be analogous to the justification in the personal case.

Rodin discusses this notion of national-defence at length in the second half of his book, ultimately rejecting it.⁸² I will not discuss his arguments here as they treat the question of *ad bellum* justification rather than *in bello* just conduct. One question concerning national-defence is relevant to non-combatant immunity, however, and this I will consider. Let us assume for the sake of argument that an analogy between national-defence and self defence can be successfully constructed. How will this explain why it is that only the enemy's soldiers may be killed in an act of national-defence, and not the enemy's civilians?

⁸¹ It is true that two enemy soldiers may unexpectedly stumble across one another in a lonely spot. In such a case, given that each may reasonably expect the other to kill him, we might accept that the situation is roughly analogous to the personal case of self defence. But these are by no means the circumstances in which most killing of soldiers takes place. Also, note that in these agent-to-agent cases where soldiers kill in self defence, there is no principled reason why soldiers and civilians would be separated. It is perfectly possible that a soldier stumble across an armed civilian and exercise his right to self defence.

⁸² See in particular chapter 6, p. 122ff. On p.125 he provides a taxonomy of the different possible combinations of the subjects, objects, and ends of the right to national-defence, and bases his following discussion upon this taxonomy. He concludes that no combination of subject, object and end can ground a right to self defence which permits the kind of killing which self defence in war is usually thought to permit.

It is worth noting that even if an analogy can successfully be constructed between national-defence and personal self defence, there is still a problem in justifying the acts of those soldiers whose state possesses no right to national-defence. If an individual soldier's attacks are only justified or rendered proportionate by the permission to kill in self defence held by the state for which he fights, then his permission to kill other soldiers is dependent on that state's permission to defend itself. But whether a state has a permission to defend itself against a given attack is an *ad bellum* and not an *in bello* consideration, and moreover, one which does nothing to explain the permissions and restrictions allotted to the unjust army's soldiers. Thus *in bello* questions of rightness are reduced to *ad bellum* ones, and as I argued in section 3 of chapter 1, it is important to keep these two sets of policies distinct.

It seems that the most promising avenue of explanation here is the requirement of necessity. It is surely necessary, in order to defend one's nation against an aggressor, to target and kill that aggressor's soldiers, but not necessary to target and kill its civilians. According to self defence, then, as long as the soldiers' deaths are in proportion to the good being defended, the killings will be justified. (Thus in the case of national-defence, the proportionality requirement of *jus in bello* corresponds to the proportionality requirement of self defence.)

This amounts to a 'military necessity' justification for the targeting and killing of combatants, and I return to it in my final chapter. I believe it is more promising than any other justification for the principle of non-combatant immunity that we have yet seen; but I also think that the consequentialist considerations at its base can and should be applied differently within *jus in bello*.

Necessity

I now turn to another application of necessity within self defence. In the personal case, killing in self defence is justified only if it is necessary. This means that the threat being deflected must not be readily deflectable in any way other than by the use of lethal force. As I mentioned in chapter 3, there are interesting questions about what degree of risk one is required to take upon oneself in order to avoid having to use lethal force against an aggressor. However, for the purposes of argument, I will assume the permission is greater rather than more constrained, and that we may kill to protect ourselves whenever not doing so would require that we take any non-negligible risk upon ourselves. What will the outcome be of such a minimal constraint for soldiers killing in war?

On one level, it will certainly render most killing in war unjustified by self defence. It is nearly always the case that a soldier may avoid a direct threat to his life by retreating, rather than by killing his attacker. This is particularly the case when we consider pilots bombing from the air, or the launching of long-range missiles: these kinds of attacks are made when the attacker's own life is not in danger at all (or at least not from those he is directly attacking). Clearly at an agent-to-agent level, the necessity constraint will not allow anything like our currently permissible range of military activities. On the other hand, let us again consider the notion that those

launching long-range missiles and pilots bombing from the air, as well as soldiers on the ground, are involved in a collaborative act of national-defence. As I mentioned above, it is surely the case that killing the enemy's soldiers is necessary to deflect the threat posed by the enemy as a whole. If we view the necessity constraint on this level, then the self defence requirement of necessity will equate to the *jus in bello* requirement of military necessity: An act of war is justified if it is necessary to the successful conduct of the war. Again, I return to this notion in my concluding chapter.

I have argued that it is not in fact an analogy between the personal case of self defence and the kinds of killing which happen in war will be successful as the requirements of proportionality and necessity become nonsensical. I have also suggested that a better way to understand a soldier's killing might be as part of an act of national-defence, and I consider the implications of this further in my final chapter. Now I will briefly talk about the problems self defence justifications face in identifying the aggressor in a war situation, before examining the particular problems which arise for self defence *even if* we assume that soldiers really may be killed in self defence.

Identifying the aggressor in war

One of the most basic problems with using self defence to justify acts of war is that there are important, and perhaps often unanswerable, questions about who has initiated a given conflict, and this is crucial to the success of a self defence justification. When both parties to a conflict claim to be performing specific acts of violence against one another *in self defence*, at least one (and possibly both) of them must be wrong in their claim.

It is possible that both parties begin to aggress against one another simultaneously, but I am doubtful as to whether this means they are both acting in self defence. It seems more reasonable to say that they are both aggressors. Intrinsic to the idea of self defence is that it is a response to a threat, and if (for example) two people start

throwing punches at one another at the same moment, neither is (at least at first) responding to the threat of the other. Rather, they are both acting aggressively.⁸³

Alternatively, one might claim that, once the fighting has begun, each side is clearly acting to defend itself against the acts of violence perpetrated by the other, and therefore both sides *can* in fact both claim to be acting in self defence. However, in the individual case we have seen that self defence is justified only when the victim is not culpable for the attack inflicted on him. Thus the question of which party to the conflict is culpable for starting the conflict is highly relevant to whether defensive permissions may be appealed to at all.

I cannot offer a solution to the problem of which side, if any, can legitimately claim to be acting in self defence in the case of a complex modern violent conflict. Indeed, I cannot imagine how one might begin to determine who holds responsibility for starting the conflict. All I wish to do here is point to a basic problem for arguments from self defence: at most one side can legitimately appeal to these arguments. Therefore, if self defence upholds the combatant/non-combatant distinction because it justifies attacking the enemy's army, it only does so for one side of the conflict. As I explored in section 3 of chapter 1, however, the real-world effect of this may well be negligible as both parties to any conflict will always claim to be acting in the right (though it may be somewhat more difficult for an aggressor to claim to be engaged in a just war *of defence*, rather than simply a just war.)

Section 2: Defending the principle of non-combatant immunity via self defence

The introduction of potentiality

The development of modern weaponry has damaged the credibility of a self defence justification for killing in war. The historical appeal of the justification was due to way battles were fought: Two forces stood ready at either side of an area which was

⁸³ One could argue that they are both responding to the potential threat posed by the other, and both seeking to remove that threat by acting more violently than their opponent. I discuss the problems with this idea of potentiality below.

to become the battlefield, and proceeded to charge at one another and try to kill as many of the enemy as possible. Battles were very much like large clusters of hand-to-hand combats, happening in the same place and at the same time. But when long-range weapons, including firearms, missiles, aerial bombs, and landmines, were developed, war began to be waged long-distance, and the henceforth imminent threat posed by the enemy evolved into a variety of direct and indirect threats open to a variety of interpretations. No longer was someone placing you in mortal danger only if they were coming at you with a sword; they could be equally dangerous hundred of miles away at a communications tower. Self defence in war changed from a deflection of an immediate threat, as present in cases of personal self defence, to something more concerned with potentiality.

The nature of modern warfare dictates that arguments from self defence in war must point to the potential, as well as the imminent, threat posed by the enemy. This is not at first glance problematic, as personal self defence also allows us to defend against what are strictly speaking only potential threats, so long as we have a reasonable level of certainty that they pose an imminent threat. For example, just as I can shoot the person shooting at me or my friends, I can also shoot the person bending to pick up a gun who I reasonably believe is intending, in the very near future, to use it to shoot at us. However, there are some merely potential threats against whom we may not defend, for example, the person who we know intends to kill us and who will have access to a gun in a week's time. The natural reason to offer for this is that we cannot have the required level of certainty that our killing in such a case would be a truly necessary act of self defence. When the threat is not imminent, there seem many more desirable ways to avoid it than by the infliction of lethal force. Likewise, in war, it is intuitive to say that the youth who will, in another six months, be old enough for the draft is not an immediate enough a threat to justify killing him in self defence. Here again, I believe our intuition is that there are probably other tactics which we could better employ to deflect the threat. I return to this idea below.

In the personal case and in war, then, there must be some degree of immediacy, some standard of reasonableness for expecting that someone will pose a future threat, which will determine whether a killing can be justified by an appeal to self defence. I do not intend to argue here for any specific measurement of threat, either for the

personal case or for the application of self defence to war. Rather, I argue that *whatever* this degree of immediacy is, if it encompasses military targets, as the combatant/non-combatant distinction requires, then it must also encompass many civilian targets. Thus if arguments from self defence succeed in endorsing acts of war against combatants, they must also endorse some attacks on those we usually class as non-combatants. I conclude that as a justification for acts of war, self defence does not divide the targets for attack into combatants and non-combatants, and therefore cannot be appealed to as the basis for *jus in bello*'s distinction between them.

Legitimate attacks against potential threats

There are many acts of war against the merely potential threat of the enemy's military which are considered legitimate. A good example is the bombing of militarised zones from the air. In many cases, a militarised zone does not represent a front against the opposing army, so people in the zone are not actively engaged, right at that moment, in killing the enemy. They do not pose an immediate threat to the lives of any member of the enemy army. However, I argue that they might be thought to pose a potential threat in two ways: first, because they may later be engaged in direct warfare; and second, because they are acting in a military capacity in support of those currently engaged in killing the enemy. It seems to me that either element of this dual potentiality might be considered sufficient to justify their targeting them. The first relies on a notion of pre-emptive self defence, and the second, indirect self defence. I explain and examine each of these below.

Indirect self defence

As a legitimate act of warfare, the only limit placed on the accepted combatant casualties of bombing a militarised zone is *jus in bello*'s general requirement of proportionality. There need be no special demonstration of the necessity of the combatants' deaths, and no indication that measures have been taken to minimise loss of life; in fact, arguably, the more killed the better (within the limits of proportionality), as far as the acting state is concerned. These combatant casualties often include a wide range of non-fighters, including mechanics and administrative and medical staff, who are classed as combatants because of their direct involvement in the war in military areas. If their deaths are justified by an appeal to self defence,

this must be because the threat they pose (or help others to pose) is regarded as sufficiently immediate, though indirect.

Suppose we grant that the threat posed by a mechanic working on an armoured vehicle in a militarised zone is sufficient to make him a legitimate target in time of war. It is hard to see how that threat is importantly more immediate than, for example, that of a mechanic working on similar vehicles preparing to drive into the militarised zone. If such a person is a threat, then so presumably are the people transporting the vehicles from the factory to the military areas. If they are a threat, so are the people assembling the vehicles at the factory... and so on. The difference of degree of threat between each class of people (and bear in mind that these classes can be much more narrowly distinguished than I have outlined here) is so slight that to say one class is a threat and the other not seems unwarrantable. However, we end in deciding that people who are usually classed as civilians pose a threat to the enemy army and can therefore be targeted. This is precisely what the immunity theorist was trying to avoid.

The problem with targeting those acting in support of the actual aggressors, and thus with indirect self defence, is that there is no reason to think that official, front-line, uniformed (and therefore combatant) supporters of the war do any more harm, or pose any more of a threat, than all the non-uniformed civilians also acting in support of the war. In fact, the reverse is often the case. Compare a military base in which encrypted messages are received to the civilian base from which they are sent. In such a situation it seems that the threat is, if anything, more immediate, or at least more serious and dangerous, at its source, in the civilian base. It would therefore be natural to suppose that any appeal to self defence would deem it more justified to bomb the civilian base than the military one, because this is a more effective way of removing the threat. It is the civilians who are, here, posing a greater indirect threat to the enemy.

This same consideration may be used to defend the direct targeting of many groups of civilians in times of war. One could argue that civilians working in factories supplying the army, or working on communication or supply lines, are preparing or aiding an attack to the same extent as soldiers in uniform at the other end of the

communication or supply lines, receiving the messages or unloading the goods. So if members of the latter group are legitimate targets because of concerns of self defence (as we usually think), then the former must be too. Indeed, when we consider that attacks on the civilian end of the lines can often be carried out at much less risk to the perpetrator of the attack, and no quantitative difference in loss of life, it seems to be a much better idea.

I do not wish to offer support to any argument whose conclusion is that warfare on civilians is a good thing. However, there is undeniably something to be said for the factual claim the argument makes regarding the parity of threat posed by civilians and soldiers. It does serve to illustrate that an appeal to indirect self defence cannot be used to distinguish between civilian and combatant targets. The problem with the argument I have sketched above is that the conclusion it arrives at, that we should be less concerned about targeting civilians, is the opposite to the one I think we should favour, that we should be more concerned about targeting soldiers. As I argue in my final chapter, our presumption ought to be strongly in favour of limiting legitimate targets and protecting life.

Pre-emptive self defence

As well as targeting those posing an indirect threat, it might be thought acceptable to target those who will be posing a threat, but are not at the present time. This includes soldiers who are moving from place to place, or who are asleep, or eating, or engaged in administrative tasks. That they are not currently engaged in shooting at the enemy does not exempt them from being targets.

The extent to which an attack may be pre-emptive is a hard question, and one I consider carefully below. First, there is the more basic question of whether an attack may be pre-emptive at all.⁸⁴ The police are not permitted to arrest people on the

⁸⁴The question of whether *ad bellum* pre-emptive strikes are permitted came to the fore when the USA invaded Afghanistan (2001) and Iraq (2003) with the alleged purpose of disarming a potential threat. Although legal viewpoints varied greatly, most commentators seemed to agree that there are at least some circumstances in which a pre-emptive strike was morally permitted. The difficulty lies in settling upon a satisfactory level of, and standard of proof of, the threat being pre-empted. Walzer gives some discussion of this question of 'sufficient threat' in *Just and Unjust Wars*, p. 80-5. Another interesting article is Neta C. Crawford's paper 'The Slippery Slope to Preventive War' (part of *Ethics and International Affairs* 'Roundtable on Evaluating the Preemptive Use of Force', vol. 17, no. 1, 2003).

grounds that they have the potential to commit crimes; does this not give us reason to doubt that states are permitted to launch pre-emptive strikes against idling soldiers on the grounds that they have the potential to pose a threat?

The comparison is not a good one, and in fact an appeal to common law supports the idea that some kinds of pre-emptive defensive strikes are not only permissible but also desirable. When, for example, the police have a wealth of strong evidence indicating that a particular crime is planned by a particular individual, then it seems that this ought to justify an arrest, particularly if the criminal has a long history of committing similar crimes. Moreover, we consider it desirable that the police act on this kind of advance information, rather than waiting for the crime to occur. Now in the case of war, there is plenty of strong evidence to show that a sleeping soldier is highly likely to be posing a threat to the enemy in the very near future, particularly as he has been doing so consistently in the recent past. Therefore the cases in which pre-emptive self defence is appealed to in war are analogous to common-law cases in which such a justification would be acceptable and reasonable. As to those cases in which, despite strong evidence of a future threat, we do not consider a pre-emptive strike justified, this is likely to be because we judge that the threat could be otherwise deflected. This is probably not going to be the case for pre-emptive self defence in wartime unless it is expected that the war can be ended before the sleeping soldier awakes.

The central problems encountered when appealing to pre-emptive self defence in wartime, just as with indirect self defence, have to do with degree: Just how pre-emptive is self defence allowed to be? If we can attack a frontline soldier who is asleep, then can we attack a soldier who is being treated in hospital? What about a soldier in training? What about a soldier on leave? What about a young man who is about to be drafted? And so on. Let us consider the first and last of these cases. It is intuitively and legally⁸⁵ acceptable to target a soldier while he is asleep, because he will be a threat when he wakes up. It is intuitively and legally unacceptable to target a civilian who is eligible for the draft in another few days merely because he will be a

Crawford provides a clear discussion of the central difficulties concerning international law and stipulating a threshold for the permissible use of pre-emptive force.

⁸⁵ I refer to such laws for just conduct in warfare as are specified by the Geneva Conventions and Protocols and by the 'rules' of the Just War tradition.

threat when he joins up. What are the relevant differences here? First, there is a difference in time period: it will be at least a few weeks before the young man poses a threat to the enemy, whereas the soldier will be a threat soon after he wakes. But remember that there are times when it is a certainty that a conflict will be a protracted one. If the enemy can be *certain* that the young to-be-draftee will in fact be a threat in the future, why can he not be targeted? Justification for a pre-emptive strike cannot seem to ask for anything more than the *certainty* that the target is going to be a threat. If we have no expectation that the threat can be deflected without recourse to lethal force, then it seems there is much less reason to require that the threat is imminent.

Another difference between the sleeping soldier and the draftee is that the soldier has been a threat already, and will be resuming the status of threat, whereas the young man has not yet done anything to make him a threat. This, however, is not a relevant concern to an act of pre-emptive self defence, but rather some sort of appeal to innocence, a type of argument we have seen to fail. The past history of the target will only be relevant in determining how likely he is to be a threat again, i.e. how likely it is that a strike is indeed pre-emptive and not simply aggressive. But as we have already granted that the young man is certainly going to be a threat, this consideration is not important.

For these reasons an appeal to pre-emptive self defence, like an appeal to indirect self defence, fails to distinguish between combatants and non-combatants in the way the immunity theorist requires. I conclude that any appeal to considerations of self defence to support the distinction is illegitimate. I now consider the possibility of using the notion of self defence to sharpen the vague terms ‘combatant’ and ‘non-combatant’ in order to prevent some of the arguments I have given here hitting their target. I again conclude that the appeal fails to deliver what the immunity theorist requires.

Section 3: The paradigm combatant as a non-innocent threat

When we think of the paradigm combatant, we think of someone who is both non-innocent (in some moral sense relevant to his involvement in the war), and who poses

a threat to the enemy. The paradigm non-combatant is correspondingly innocent of involvement in the war, and does not pose any threat to the enemy. The appeal and justification of the principle of non-combatant immunity rest on these moral differences. In the preceding chapters I have examined these two paradigmatic qualities, arguing for the conclusion that they do not in fact identify combatants from non-combatants in the way the immunity theorist supposes. I have also explained why redefining ‘innocent’ so it equates to ‘not posing a threat’ is not a desirable move.⁸⁶ We now face three possibilities to repair the moral error contained in the principle of non-combatant immunity:

1. Alter the principle so that instead of distinguishing combatants from non-combatants, it distinguishes the guilty from the innocent, or the threatening from the non-threatening
2. Try to defend the principle of non-combatant immunity in a way which does not rely on the existence of intrinsic differences between civilians and soldiers
3. Discard the principle entirely and reformulate *jus in bello* without it

I will argue in later chapters that a combination of the second and third of these options is the most desirable choice. Here I wish to consider the first, and argue that this cannot offer the moral robustness we should require of a principle of *jus in bello*.

We have identified innocence and guilt as moral notions relevant to the conduct of war. Likewise, constituting a threat seems to affect one’s moral status as a target, due to others’ rights to self defence. The problem for the immunity theorist is this: the three pairs of categories we have identified (innocent/non-innocent, threat/harmless, combatant/non-combatant) by no means overlap neatly in the way that immunity theory implies. It is not the case that all the non-innocent pose a threat, nor that all the innocent are harmless to the enemy. Nor is it the case that all combatants are non-innocent and threatening, and all civilians innocent and non-threatening. Therefore if we want to argue that one of these features really is morally important and deserves to be recognised in Just War theory, we must choose which quality (or which

⁸⁶ See chapter 2, section 1.

combination of qualities) is sufficient to make someone a legitimate target, and which are sufficient to render their possessors exempt from attack.

The immunity theorist might suggest that we accept the class distinction between combatants and non-combatants as the most morally stringent, and therefore base immunity theory on the presumption that simply being a combatant is, of itself, what makes combatants liable to be targeted. But this claim is morally unsatisfying, because it begs the vital question of *why* combatants are morally vulnerable to attack. If the immunity theorist's answer is just 'Because they are combatants, and all combatants are legitimate targets', then we will naturally reply that this assertion requires some argument. Then the immunity theorist must either go back to square one, and argue for the proposition that all combatants are legitimate targets by pointing to some morally relevant quality possessed by all combatants; or she must argue that her claim is morally foundational. I have indicated the problems with the former possibility; the latter is entirely without intuitive plausibility. We may well have the intuition that combatants are legitimate targets of attack and non-combatants are not, but we all also certainly have the intuition that this has something to do with the qualities possessed, and the activities performed, by these two groups. The categories of combatant and non-combatant as we define them are artificial, and there is no reason to suppose them morally basic.

A more promising possibility for the immunity theorist is to decree that the possession of some morally relevant quality is what really matters to retain one's invulnerability to attack. The immunity theorist could point out that as well as being synthetic, 'combatant' and 'non-combatant' are vague terms. She might suggest that we use some morally relevant quality, like innocence or the constitution of a threat, to determine where the line between the two groups ought to be drawn. Unlike the previous position, this has the advantage that we do generally agree that innocence and the possession of a right to self defence are important moral qualities which ought to feature in debates about the protection of life; and we also agree that finding a morally relevant difference between two groups justifies treating them differently.

If we decide that innocence is what matters, then, we will accordingly judge that combatants ought to be properly defined as those who are culpably responsible (to

some fixed degree) for any acts of war, and non-combatants as those who are not so responsible. If we decide that the constitution of a threat is a more important moral feature, then anyone who poses a fixed degree of threat will be a combatant, and all those who do not will be non-combatants. Although each of these methods is more promising than the outright claim that the combatant/non-combatant distinction is a morally foundational one, I argue here that they do not offer acceptable solutions to the immunity theorist's problem of definition. I will first consider the case of innocence, then of constitution of a threat.

Innocence as a criterion for combatant-hood

Suppose we agree that what distinguishes combatants from non-combatants is the innocence or guilt of the agent. If the immunity theorist wants to preserve the basic form of non-combatant immunity, she will now argue for her principle this way:

1. All combatants are non-innocents (as per new stipulation)
2. All non-innocents are legitimate targets
3. Therefore, all combatants are legitimate targets.

4. All non-combatants are innocents (as per new stipulation)
5. No innocent is a legitimate target
6. Therefore, no non-combatant is a legitimate target

Thus the immunity theorist retains her conclusion, that combatants (now defined as non-innocents) may be targeted and non-combatants (now defined as innocents) may not; and she can provide a moral basis for this claim.

One immediate problem for the immunity theorist is that even if we suppose that we *could* describe exactly what it meant to be non-innocent to the degree required to render one a legitimate target of attack, it seems highly unlikely that modern warfare is going to allow us to identify these people from, for example, a bomber plane, or a long-range missile launch-pad. Indeed, it is open to debate whether country A will *ever* have the means to decide which of country B's members are morally non-innocent, or morally responsible for their involvement in the war, whether they are bombing from the air or sniping across a field. To judge a soldier's moral guilt,

remember, we must not only judge the guilt of his actions, but demonstrate that he is morally responsible for those actions. I have argued in chapter 2 that where conscription operates with some penalty attached for refusal, there is good reason to think that many soldiers in the opposing army are not morally responsible for their position; or at least, not sufficiently to render them targets for potentially fatal attacks. This is just a rephrasing of the original problem of arguments from innocence: not only does 'non-innocent' not map onto our common conception of 'combatant', but it is hard to see how *any* genuinely moral definition of innocence could ever identify a sufficiently geographically homogenous group which it would be morally legitimate to target. If the immunity theorist wants to redefine 'combatant' to mean 'non-innocent', she may end up with the contingently pacifist conclusion that although combatants are, in and of themselves, legitimate targets of attack, the nature of their dispersal in the population means it will never be legitimate to attack them. The only exception would be cases where the enemy knows the soldiers they are targeting to be willing, well-informed fighters freely embarked on furthering an unjust cause. I cannot imagine a modern war that fits this characterisation; in any case, the immunity theorist does not wish her theory to apply only in such cases.

Constitution of a threat as a criterion for combatant-hood

Will the immunity theorist fare better if she chooses to use the constitution of a stipulated degree of threat as her criterion for distinguishing combatants from non-combatants? This would be a way of avoiding the charge I discussed earlier, that it is wrong to use arguments from self defence to defend the practice of targeting combatants because they do not all pose a like degree of threat.

First, as with the innocence criterion, there will be much difficulty in choosing just what degree of threat, and with what causal directness, and over what time period, is necessary or sufficient to render the possessor liable to attack. However, unlike the innocence criterion, it is probable that those posing a like degree of threat will be geographically concentrated together. For example, all soldiers in a hospital and medics attending them will pose a similar degree of threat to the enemy (though possibly different degrees of long-term threat), and it is desirable (both in moral and practical terms) that they are judged all non-combatants. All soldiers at the front-line of fighting pose a similar threat, and it is desirable that they are all judged

combatants. All workers in munitions factories pose a similar degree of threat, and whether they are judged combatants or not, it is at least *prima facie* desirable they are all judged alike.

So let us imagine that the immunity theorist can calculate what degree of immediate and/or indirect and/or long-term threat someone must pose to the enemy in order to render themselves liable for attack. All those who fall on or above this threshold will be considered combatants, and all those who fall short of it will be considered non-combatants. All combatants will be legitimate targets, and no non-combatant will be a legitimate target. The problem is this: In all other ways than that of ad-hoc definition, *this would not be non-combatant immunity at all*. (Indeed, this kind of threat-related moral distinction between legitimate and non-legitimate targets is similar to one I will argue for in the later part of this essay, but I would not consider myself an immunity theorist.) If we judge that soldiers at rest or engaged in administrative tasks are not combatants, but ‘civilians’ working in communications towers or designing weaponry on computers in inner-city office blocks are, then we make nonsense of the intuitively agreeable idea which immunity theory was meant to encapsulate.

A problem of ad-hoc definition

To alter the definition of combatant so markedly as to incorporate into it all threats, and that of non-combatant so it excluded all threats, is to render the choice of the labels ‘combatant’ and ‘non-combatant’ for these categories arbitrary, in that the commonsense interpretation of the label does not fit the content. Likewise, to judge all non-innocent soldiers to be *by definition* combatants and all innocent soldiers to be *by definition* non-combatants, regardless of other features of their activities, is an inappropriate use of moral language.

The word ‘combatant’ is not arbitrary in this way; it is clearly intended and understood to indicate those engaged in warfare, and with this understanding is the intuition (justified or not) that all combatants are non-innocent in some morally relevant way.⁸⁷ Thus I state again: The principle of non-combatant immunity states that combatants may be attacked in war and that non-combatants may not. The

⁸⁷ Although the term is a vague one, to sharpen it in either of the ways I have suggested above is to alter its extension too markedly for it to retain any meaning.

paradigm combatant both is non-innocent (in some sense relevant to his involvement in the war), and poses a threat to the enemy. The paradigm non-combatant is innocent of involvement in the war, and does not pose a threat to the enemy. If the immunity theorist chooses either of these moral notions and restricts the application of 'combatant' and 'non-combatant' to those to which this notion applies, she not only runs into practical problems to do with definition and geographical dispersion, but she changes her theory from non-combatant immunity to (something like) innocent-immunity or non-threat-immunity. It is perhaps acceptable, in searching for a morally sound definition, to sacrifice some of our intuitions about who is and who is not a combatant, as it would seem these terms are indefensible as our intuitions cast them; but to redefine the terms so that they do not encapsulate their paradigmatic characterisations at all is to alter the theory at the cost of its very nature and intention.

Let us revisit the three possibilities for the immunity theorist I identified above:

1. Alter the principle so that instead of distinguishing combatants from non-combatants, it distinguishes the guilty from the innocent, or the threatening from the non-threatening
2. Try to defend the principle of non-combatant immunity in a way which does not rely on their being intrinsic differences between civilians and soldiers
3. Discard the principle entirely and reformulate *jus in bello* without it

Of these three possibilities, the above discussion explains why the first is not an option for the immunity theorist who wants to preserve her theory. We are left with a choice between defending the principle of non-combatant immunity by referring to extrinsic characteristics of civilians and soldiers, as a utilitarian might; or discarding it entirely and reformulating *jus in bello* without it. I will spend the next chapter considering this first option, and the final chapter considering the plausibility of the latter option in the long term.

Section 4: Conclusions as regards immunity theory

I argued in chapter 2 that if arguments from innocence really do serve to protect civilians from attack on the basis of their guilt, then the immunity theorist must explain why it is that innocent soldiers are not protected. Similarly, if it is the guilt of soldiers which allows them to be targeted, then we need an explanation of why guilty civilians are not liable to be likewise targeted. The immunity theorist must either answer each of these charges, or accept that moral innocence doesn't make a difference to the legitimacy of targets, in which case an alternative argument must be offered to support the distinction between civilians and soldiers. I suggested that such an alternative might be an appeal to self defence, the justification of which I explored in chapter 3.

My arguments in this chapter have been procedurally very similar to those in chapter 2. I have claimed that if arguments from self defence are successful in justifying the targeting of soldiers, then the immunity theorist is in trouble, because these arguments also show that it is legitimate to target the civilian population in a large number of cases. If the arguments are not successful, however, then the immunity theorist is still in trouble, because she must offer an alternative justification for the combatant/non-combatant distinction.

I know of no other justification, outside of divine decree, offered for the claim that combatants and non-combatants are intrinsically morally different in such a way that the killing of one group in war is wrong, but the killing of the other group is not wrong. In the absence of further argument, then, I arrive at the following conclusion: Any reason offered for why it is generally wrong to kill, or why it is sometimes acceptable to kill, does not indicate the existence of any intrinsic difference between civilians and soldiers. It is true that there are clear cases in each class of persons, but likewise, there are clear exceptions.

I conclude that the principle of non-combatant immunity is not an intrinsically moral principle and does not point to a real moral difference between civilians and soldiers. If it is to be justified as part of *jus in bello*, therefore, that justification must be found extrinsically, in the role the principle plays in generating desirable consequences. I examine this possibility in the next chapter.

CHAPTER FIVE

Policy and moral conventions

I have argued that there are good grounds for thinking that the principle of non-combatant immunity is not founded on an intrinsic moral distinction between civilians and soldiers. However, there remains the possibility that its place in *jus in bello* can be justified with reference to extrinsic facts about the role it plays in governing warfare. In this chapter I examine this possibility.

I first look at an argument offered by George Mavrodes for the moral value of non-combatant immunity in the absence of any intrinsic moral distinction between soldiers and civilians. His argument states that a moral obligation to follow non-combatant immunity results from the wide acceptance of the convention. I show that this argument implicitly relies on a utilitarian justification: it reasons that we ought to follow the principle of non-combatant immunity because it creates a better outcome than not following the principle. Judith Lichtenberg uses this reasoning explicitly in her 1994 paper ‘War, Innocence, and the Doctrine of Double Effect’⁸⁸, and in section 2 I discuss the wider consequences of justifying *jus in bello* with this type of argument. Then in section 3 I identify the two tasks to which political philosophers might attend, and the significance of this to the arguments offered by Mavrodes, Lichtenberg and others.

Section 1: Convention-dependent morality and utilitarian heuristics

⁸⁸ Lichtenberg, Judith, ‘War, Innocence, and the Doctrine of Double Effect’. *Philosophical Studies* 74, 1994, pp. 347-368.

George Mavrodes, in 'Conventions and the Morality of War', argues that although attempts to find intrinsic moral differences between civilians and soldiers fail, the principle of non-combatant immunity can nevertheless be defended because it is a widely accepted convention.⁸⁹ He argues that some conventions create moral obligations which would not exist if the convention did not exist, and gives traffic laws, and our obligation to follow them, as an example:

It does not seem plausible to suppose...that one can discuss these duties without immediately referring to our laws and customs. And it seems likely that different laws would have generated different moral duties...One cannot account for the specific obligation apart from the convention.⁹⁰

Thus I am not morally bound to drive on the left because of any quality either of mine or of the left side of the road. But because it is a convention in New Zealand to drive on the left, I have a responsibility, a moral obligation, to keep to this convention. Mavrodes believes that the principle of non-combatant immunity is, like traffic laws, a convention rather than an intrinsically moral principle. However, as with traffic laws, the widespread acceptance of this convention creates a genuine moral obligation, and therefore the principle has moral weight.

Will any convention do?

There is something missing from Mavrodes' account of convention-dependent morality as it is presented here. While it is true that some moral obligations exist only because of the conventions or customs or laws behind them, it does not follow that *any* convention creates a moral obligation, even if widely followed. If, for example, our traffic laws were grossly discriminatory in a way which caused widespread suffering, they might not create a genuine or lasting moral obligation even if they were widely accepted. Likewise, just because the convention of non-combatant immunity is (or at least claimed to be) widely adhered to, that is not enough to say it

⁸⁹ Mavrodes, G., 'Conventions and the Morality of War', in *Philosophy and Public Affairs* 1975: vol. 4 no. 2, pp. 117-131.

⁹⁰ Mavrodes, p.126-7

creates a true moral obligation. Wide acceptance may be a necessary condition, but it is not sufficient. For a convention to succeed in creating a moral obligation, it must possess something more than wide acceptance.

I propose that this ‘something more’ is that following the convention must serve to avoid inflicting harms and promote the attainment of goods; it is a consequentialist consideration.⁹¹ The convention of driving on the left in New Zealand creates a moral obligation to do so for the same reason which guarantees its wide acceptance: it makes driving safer and easier for everyone, thus creating a desirable situation. However, if there were a convention which said that people over 6 feet tall could travel during the day but shorter people could only travel during hours of darkness, that convention would not create a moral obligation, no matter how widely accepted it was. This is because breaking such a convention would better serve the good than adhering to it. Both principles make intrinsically arbitrary distinctions, but one is arbitrary in a harmful way, the other in a beneficial way.

(It is tempting to say that one principle is inherently discriminatory and the other is not. But this is not so; in fact, the enforcement of either principle would discriminate against some group, whether that be people under 6 feet tall, or people who prefer to drive on the right hand side of the road. Neither principle rests on an intrinsically moral distinction; in both cases the distinction is, as in the case of non-combatant immunity, non-moral. For the distinction to be rendered immoral (rather than just non-moral), it must be extrinsically so, i.e. adoption of some particular convention utilising the distinction must create an undesirable outcome. Note that not all conventions based on a given distinction will render it immoral. There is nothing immoral about distinguishing between people because of their height when, for example, deciding admittance to rollercoaster rides.)

I argue that a convention succeeds in creating a genuine, true and lasting moral obligation only if it meets two criteria:

⁹¹ The moral justification for conventions of the sort Mavrodes discusses must be of a consequentialist or otherwise non-deontological nature, and this is why: Any justification referring to the intrinsic rightness of the convention begs the question, which is asking which conventions create a moral obligation *which would not exist if the convention did not exist*. If the convention in question is an expression of some sort of deontological moral requirement, then it is not the same sort of convention as non-combatant immunity.

1. Adhering to the convention will, in general, create a better moral outcome than breaking the convention;⁹²
2. There is no other convention available which will also be widely accepted and which better serves the moral good.

Below I assess Mavrodes' position by considering the principle of non-combatant immunity with reference to these criteria. The principle will meet the first criterion if it can be shown that wars in which only soldiers are targeted cause less harm than wars in which anyone may be targeted; I argue for this conclusion in the following section. The principle will meet the second criterion if there were no other generally acceptable principle which would deliver an even better result than non-combatant immunity. I consider this possibility now.

The possibility of other conventions

Mavrodes recognises that, when we assess a convention, we must consider the alternatives to that convention. If there is a better convention available to be followed, then that better convention ought to be adopted.⁹³ If, for example, it could be demonstrated that traffic accidents would be significantly less frequent if everyone in New Zealand drove on the right instead of the left, then we (as policy makers, conscientious drivers and/or concerned citizens) would have a moral obligation to

⁹² This first criterion is very similar to what a rule utilitarian might offer as a justification for accepting a certain rule as a moral one. A rule ought to be followed if following it will, in general, create a better outcome than breaking it. The problem for the rule utilitarian is in explaining why the rule ought to be kept to in cases where secretly breaking it seems to attain a better outcome, as in the classic 'lonesome stranger' problem. The same problem might arise for conventions: if there is an occasion on which violating non-combatant immunity will generate a better outcome than adhering to it, does there remain a reason to adhere to it nonetheless? I would answer in the negative; but I would add the caveat that, as I discussed with reference to 'rules of thumb' in section 3 of chapter 3, the necessary foreknowledge is so hard to come by, and the utility calculation so difficult, that the question is unlikely to arise in the real world.

I do not deny that there are real-world hard cases about direct attacks on civilians; the atomic bombs dropped on Nagasaki and Hiroshima represent such a case. Some historians and military theorists claim that despite the massive loss of life, not only more people but more *civilians* would have been killed by conventional war on Japan. (For a good overview of the debate, see Bryan McNulty's article 'The Great Atomic Bomb Debate', published in the Spring/Summer 1997 edition of Ohio University's online publication 'Perspectives'. Retrieved from www.ohiou.edu/perspectives/9701/bomb2.htm.)

⁹³ Mavrodes, p.127-8

alter our policy so as to enforce the new convention.⁹⁴ However, this is not to say that any moral obligation that we conscientious drivers currently have to drive on the left would automatically vanish. Instead, we would have a real obligation to keep to this convention while progress was being made toward desirable political change.

It is important to note that the reason why we, as conscientious drivers, ought not start to follow the new, better convention right away: because if we do, then we fail to promote the good we are aiming for (in this case, avoiding traffic accidents) as well as if we had adhered to the current convention while working to have it officially changed. Once again this is a consequentialist justification for preferring a particular convention. If it were not the case that adhering to the convention created a better moral outcome than rejecting it, then we would not have an obligation to keep to the current convention while agitating for change. So the wrongness or rightness in adhering to sub-optimal conventions whilst they are under development depends on whether adhering to the convention in the short term promotes or undermines the good.

We find a similar standard in operation when we consider the wrongness or rightness of following unjust laws: in cases where breaking an unjust law will better promote the good than following it (even when only a small group of people break it), it seems there is a clear obligation *not* to follow the law. A good example is the refusal of, at first, very small groups of people to obey racially discriminatory laws in southern USA in the 1960s: their law-breaking was morally right because it better promoted the good than did compliance with the law. But there are other cases in which breaking even a bad law will do more harm than following it; and in such cases we have an obligation to obey the existing law while we try to get it changed. An example might be if a law were passed (in an attempt to raise a city council's profits to fund the end-of-year Christmas party) which stated that all urban household waste must be packaged in exorbitantly expensive council-branded rubbish bags, or it would not be collected from the kerbside on rubbish day. This would be a bad law: it

⁹⁴ The obligations of policy makers and of conscientious drivers naturally differ. The obligation of the policy maker is to develop, legislate and enforce the new convention in the best possible way. The obligation of the conscientious citizen is to bring it to the attention of those in charge that such a change is desirable. (It must be remembered that policy makers are also citizens, and thus the two kinds of obligations sometimes co-occur in certain individuals who hold political power.)

represents an abuse of the power allotted to the city council, and unfairly disadvantages low-income households. Thus there would be an obligation on behalf of all citizens to motivate for its change. However, if individuals were to break the law by failing to use the bags, knowing that their rubbish would therefore not be collected and would be left to rot on the footpath, they might create a much worse situation than if they followed the law whilst it was being reviewed.⁹⁵

In short, we have an obligation to follow an existing convention, even if there could be a still better convention in place, as long as this promotes the good. (We will also have some obligation to work for the existing convention to be changed to the preferable convention.) Below I examine whether the convention of non-combatant immunity meets this standard of promoting the good, or whether it ought to be broken even in the absence of an alternative proposal.

Section 2: Putting non-combatant immunity to the test

I have argued for the following claim: If there is another immediately available convention for fighting wars, the adoption of which would better promote the good than non-combatant immunity, then we are under no obligation to follow the latter. But if this is not the case, then we can conclude that there does exist an obligation to uphold the convention of non-combatant immunity *in the meantime*, even if there are better conventions which might be adopted in the future. (There remains the question of what these other conventions might be, and I briefly discuss this in my concluding chapter.) I have shown why Mavrodes' arguments imply this claim; Judith Lichtenberg, in an excellent 1994 paper, argues similarly that the value of non-combatant immunity lies in the role it plays. She grants that there is no intrinsic moral distinction between civilians and soldiers, but nevertheless considers non-combatant immunity to be a principle of value:

Where your victory is morally important enough, you are justified in killing the soldier *if necessary*. But likewise, where your victory if

⁹⁵ Of course, if attempts to get the law changed through normal processes were unsuccessful, there might come at a time at which dumping rubbish on the kerb by way of extreme protest *would* be justified, as a last resort. It will remain unacceptable only while there is a preferable alternative.

morally important enough you may be justified in killing civilians *if necessary*. . . [Non-combatant immunity] is a very useful convention, because it prevents war from infecting the whole of every life and destroying all that we value.⁹⁶

In simple terms, the reason we ought not to target civilians is not that they are in a different moral category from soldiers, but because war is less evil or unpleasant or violent or harmful overall if we do not.⁹⁷ It is this claim, and related implicit claims about the (lack of) alternatives to non-combatant immunity, which I examine below.

The need for wide acceptance

As I discuss in section 3 of this chapter, one of the biggest constraints on alternatives to non-combatant immunity which are *immediately* available⁹⁸ as policy choices is whether or not they stand a good chance of being made into international law and adhered to by warring states in the absence of an international law enforcer. Any proposed convention which will not be adhered to by any states currently at war (even to the limited extent that they adhere to non-combatant immunity) will not better promote the good than non-combatant immunity, which although far from ideal, at least places some limit on the harms of war. Of course, it is highly unlikely that any convention not agreeable to powerful states stands a chance of being made into international law at all; but even if it were, it would have to be followed in order to generate a desirable outcome. If we, as a global political community, were to abandon non-combatant immunity in favour of some policy which we knew no-one would follow, we would stand to lose the benefits of non-combatant immunity without any compensating gain by way of new protections from the harms of war.

To negate the apparent existing moral obligation to observe non-combatant immunity, then, there would need to be a convention, acceptable to states currently at war, which

⁹⁶ Lichtenberg, p. 363. Italics in original.

⁹⁷ I am assuming throughout this essay that ‘promoting the good’ and ‘avoiding inflicting harm’, as far as fighting wars is concerned, means preventing the suffering of innocents as far as possible, and in particular, minimising the loss of innocent lives. As I argued in chapter 2, many soldiers are included in this category of ‘innocents’. Therefore promoting the good and minimising harm in war involves generally limiting the casualties of war. I discuss the importance of this further in my final chapter.

⁹⁸ By ‘immediately available’ I mean available in the absence of fundamental changes to the power structures of the political world and to the way international law is made and enforced (or rather, not enforced).

would better limit such harm than non-combatant immunity. Only if such a convention or principle could be found would non-combatant immunity have no moral weight. Below I explain what I think the only immediately available alternatives are, and why they fail to meet this requirement.

Alternatives to non-combatant immunity

In the absence of international law enforcement, there are probably two main forces which constrain how states fight wars. First, their own moral sense; and second, public opinion.⁹⁹ (I am assuming that, if unrestrained by these considerations, states would simply act so as to achieve the earliest victory possible at the smallest cost to themselves.) The first consideration is, history indicates, generally not to be trusted to where, for example, public opinion is clearly in support of war, other states are not willing to act to prevent it, and economic interests are at stake. But the second consideration does often materially restrict how states behave. Therefore if we seek to identify those principles which states will pass into international law and adhere to in the absence of any meaningful form of international law enforcement, then it seems public opinion (in the broad sense described in footnote 99) will be our best motivational implement. I now consider various alternatives to non-combatant immunity and how the power of public opinion might affect their chance of acceptance.

Public opinion about how a government should conduct itself in times of war is rarely united to any degree, but one extremely widely held conviction is that targeting soldiers is acceptable and targeting civilians is not, a conviction steeped in hundreds of years of fighting tradition. This gives non-combatant immunity an advantage over many other conventions when it comes to acceptability. Because it is the status quo, it has become part of the 'moral consciousness' of warring states, and is enshrined in international law.¹⁰⁰ Violating this convention is portrayed by the media (a powerful

⁹⁹ I am including under the umbrella of 'public opinion' such diplomatic concerns as what will give credit in the international political arena, so other heads of state are thus included in the 'public'. This is because the motivation of states that I am concerned with here is that of 'keeping up appearances', or appearing to act well according to the general moral community. In general, when states appear to credit in the eyes of the public, they will also appear so in the international arena. If they are sufficiently powerful that this need not be a concern, then public opinion fails as a form of control over states' behaviour. One might think the USA now approximates this state.

¹⁰⁰ Statements about the distinction between civilians and soldiers, and the protection due to the former, can be found in Article 4 of the Third Convention of the Geneva Conventions (12 August 1949), and

influence on public opinion) as one of the most serious moral crimes of war, and for this reason states strive to be seen to adhere to it.¹⁰¹ Thus non-combatant immunity has a signal advantage over any other principle that might be offered in its place.

Consider the possible general variations on non-combatant immunity: unlimited warfare; some kind of less limited warfare; some kind of more limited warfare; or a prohibition on warfare, i.e. pacifism. As I stated above, to negate the apparent existing moral obligation to observe non-combatant immunity, there would need to be a convention, acceptable to states currently at war, which better promotes the good than non-combatant immunity. Can any of these alternatives furnish us with such a convention? I consider each below, in reverse order, with reference to this question.

The possibility of pacifism

Although there are large proportions of the general public who say that they would like to live in a world in which there was no war, I would suggest that the proportion of the general public or international community truly committed to pacifism is currently far too small to persuade governments to sacrifice their war-related interests in its favour in the near future.¹⁰² In order for governments to be influenced by their citizens' preference for non-violent means of conflict resolution, peace must seem desirable to citizens even when their national security or economic interests are threatened. In the absence of any reliable form of protection in the shape of

also in Articles 43, 44, 50 and 51 of the Geneva Protocol (additional to the Conventions of 1949) relating to the Protection of Victims of International Armed Conflicts (Protocol I, 8 June 1977).

¹⁰¹ There are few conventions of war which are similarly universally accepted and regarding which public pressure is so strong, and those I can think of tell of our concern for civilian life. First, there is a general ban on the use of landmines and biological weapons, weapons which cannot be restricted to their military targets but will always inflict harm on civilians. Second, there is much concern about the humane treatment of prisoners-of-war, as we have seen with the recent controversy over treatment of Iraqi prisoners by the US and UK occupying forces. It is interesting that the lives of enemy soldiers increase so much in value when they are forced to stop soldiering, although I think the notion of pre-emptive self defence can explain why we think it so much worse to kill a POW than a soldier asleep in his own camp. The latter alone represents a future threat.

¹⁰² Even devoted proponents of pacifism such as Robert Holmes allow that changes to the way we think about violence and victory cannot be instantaneously achieved. Holmes states that a future of non-violence requires 'a massive commitment of our moral, intellectual, and financial resources' to developing new ways of securing peace. Not only this, it demands 'a new perspective that sees the people of the world as arrayed, not basically against one another, but against the deceit, ignorance, and arrogance of governments and the ways of thinking that have produced them.' I am inclined to agree with Holmes that this is possible; but I do not think, and I am certain Holmes does not think, that it is possible in the immediate future, or without fundamental changes to the structure and content of international law. (Holmes, R., *On War and Morality* (Princeton: Princeton University Press, 1989), p. 294, p. 291.)

international law-enforcement, few people, I think, are sufficiently committed to peace to refuse to support what they view as a defensive war.¹⁰³

When it comes to a public aversion to fighting aggressive wars, the outlook is perhaps more hopeful. It is possible that a widespread civil commitment to peace could make aggressive warfare hugely unpopular or even (if the army cannot be adequately staffed with volunteers) impossible in democratic societies. Indeed, recent (December 2004) news from America and Iraq gives us reason to hope that this may be so: the United States Army National Guard has missed its recruitment target this year, the first time in a decade that this has happened; more than 5000 US soldiers have deserted the army or Marine Corps rather than go to fight in what they deem an unjust war in Iraq; 5600 soldiers from a 'national emergency pool' have been mobilised in order to supplement troop numbers in Iraq; more than half the soldiers currently in Iraq do not plan to re-enlist when their term of service expires.¹⁰⁴ It seems hopeful that a spread of this kind of anti-aggression sentiment and civil unrest – for soldiers are civilians before they become soldiers – could turn a profitable war into a political, financial and military disaster.

However, it is too quick to conclude from these encouraging signs that the current level of public commitment to pacifism is enough to prevent aggressive war. It must be recalled that resistance before the war began was not sufficient to prevent it. Although the horrors of Iraq may confer a valuable legacy in deterring future potential soldiers from joining up, we must remember that the aggressive invasion was made acceptable to many frightened Americans who were, I am sure, genuinely peace-

¹⁰³ I am acutely aware that this assertion stands undefended by hard evidence, but without being able to conduct a poll to gauge the level of support for a defensive war amongst those with pacifist convictions, such evidence will be unforthcoming. I must hope that my statement here is sufficiently intuitive to be accepted, and indeed I can offer myself as an example: Although I hold a strong conviction that war is a terrible thing, I can see that it might seem justified when compared with surrender to a hostile power who threatened violence to my person, freedom, or way of life.

We might also note that even pacifists often see a good reason for their own states not to disarm and disband their armies in the absence of everyone else agreeing to do so.

¹⁰⁴ Sources: CBS News story of December 8, 2004, retrieved from www.cbsnews.com/stories/2004/12/06/60II/main659336.shtml; and Times Online story of December 10, 2004, retrieved from www.timesonline.co.uk/printFriendly/0,,1-3-1397131-3,00.html. This latter source also states that 'Forty per cent of the 138,000 troops in Iraq are part-timers who never expected to be sent to the front line.' This fact may serve as a deterrent to those who like the idea of being in the army but dislike the idea of fighting to the death.

loving.¹⁰⁵ The aggressive invasion was made acceptable because, despite the incredulity of much of the world, it was initially viewed by many of its supporters as a defensive war. A Los Angeles Times poll of December 2002 stated that 90% of those polled believed that Iraq was developing weapons of mass destruction.¹⁰⁶ Amazingly, six months later, when Iraq had been invaded and no evidence of such weapons uncovered, 1 in 3 US citizens polled nevertheless believed that weapons of mass destruction *had* been found in Iraq; and 22% believed them to have been used against US troops.¹⁰⁷ As long as this kind of widespread misinformation can be propagated, pacifist agitators will find it difficult to commit the general population to peace, and governments will not be strongly swayed by pacifist considerations.

In summary, either too little of the Western world's population has made a true commitment to peace, or too few people consider such a commitment prudent in the absence of other kinds of protection from attack, or the governments of the nations concerned are too able to portray their wars as promoting peace, to make the pacifist position acceptable to states in today's political reality. As I argue in section 3 and in my final chapter, that should not give us cause to dismiss it as a future possibility.

The possibility of more limited warfare

Between non-combatant immunity and pacifism there are various possibilities of more limited warfare, warfare in which the targets are more restricted than under non-combatant immunity (i.e. not only must non-combatants be immune from direct attack, but also some combatants). But just as with pacifism, I think that the lack of pressure on governments to conform to more restricted targets means that such conventions stand little chance of acceptance. After all, to restrict military targeting is

¹⁰⁵ It is important here to acknowledge the protests of large sectors both of the American population, and of individuals and nations around the world, at the US-led invasion of Iraq. Certainly many people considered the war to be aggressive and unjust well before it began, and were dedicated to peace even at the cost of allowing Iraq to remain uninvaded. However, these voices were insufficient to outweigh the war-related interests of the powerful or frightened.

¹⁰⁶ Los Angeles Times story of December 17, 2002, retrieved from www.latimes.com/news/custom/timespoll/la-na-iraqpoll17dec17.story.

¹⁰⁷ Philadelphia Inquirer story of June 14, 2003, retrieved from www.philly.com/mld/inquirer/news/front/6085261.htm?1c. See also the PIPA (Programme on International Policy Attitudes) story from the same month, at www.pipa.org/whatsnew/html/new_6_04_03.html.

to make war more difficult and time-consuming and expensive.¹⁰⁸ States would need a strong inducement to adopt such a policy, and the political reality of today is that many states have no such inducement. In the absence of international law-enforcement, a rule which requires more limited targeting in war has, I think, no chance of acceptance amongst the major world powers.

Now, I do not wish to suggest that there are not individual states who are willing to adhere to more a stringent moral standard than that of non-combatant immunity. It is probable that there are; and it is certain that different nations' militaries differ in what they are comfortable with permitting by way of targeting. It is possible that some groups of nations may be persuaded to sign multilateral treaties limiting the conduct of war beyond non-combatant immunity, and this is obviously to be much desired.¹⁰⁹ Such actions could be part of a growing recognition that adhering to non-combatant immunity is not enough to guarantee the just conduct of war. However, at least in the foreseeable future, there will be a handful of hugely powerful nations who have no incentive to enter into any such treaties; they can bear the costs of fighting more easily than their enemies, they have a vested interest in warring, and (as I have discussed above) they do not stand to lose massive popular support among their voting public by going to war. When it comes to controlling the behaviour of these nations, a commitment to non-combatant immunity is, as I argue below, much better than nothing at all, and perhaps the best we can hope for in the short term.

The possibility of unlimited warfare

The chances of states openly adopting a policy of *unlimited* warfare on civilians and not being subjected to intervention or political revolution is thankfully small. Being aware of this, states are unlikely to openly endorse such a convention as part of

¹⁰⁸ This is true even if both sides were to obey the tighter restrictions, because it would mean that both sides could do less to injure the other side. The greater the extent to which this is true, the longer it will take for one side to be victorious. If, for (extreme but illustrative) example, the Allies and the Nazis were only permitted to kill enemy men in uniform actually firing at them, then World War II might still be underway.

¹⁰⁹ In fact, for these nations, such conventions as I reject for international law might be possible and thus obligatory. Bordering nations might be encouraged to organise peace treaties, perhaps strengthened by economic co-dependence, which are to the benefit of all. Retaining non-combatant immunity as part of general international law ought not be seen as a barrier to these sorts of bilateral or limited multilateral agreements. In fact, if such treaties were to become widespread, public pressure for the more powerful states to ratify them might increase.

international law. Such a convention would also run contrary to the strongest moral beliefs of the bodies which make and (to a limited extent) enforce international law.

The possibility of less limited warfare

Left remaining is some convention of less limited warfare, in which some attacks directly targeting civilians are permitted. The 20th century is littered with examples of attacks on civilians, notably the Allied bombing of Dresden and other German cities during World War 2 and the retaliatory attacks on London, and the atomic bomb attacks on Hiroshima and Nagasaki, all of which took place before the Geneva Conventions of 1949 outlawed such actions; and the US bombing of rural Vietnam and Cambodia in the 60s and 70s, which was conducted in partial secrecy and in defiance of these conventions. In both cases at least some of the actions taken have been condemned, not only by the international community but also by the successive governments of the aggressive states. It was never (to my knowledge) suggested that the convention of non-combatant immunity ought to be altered to make such actions permissible acts of war. This might lead us to conjecture that the international community would not welcome an official change to the rules of warfare allowing more attacks on civilians.

On the other hand, in the 2003 USA- and UK-led invasion of Iraq, the American troops were licensed by their government to fire not just on the Iraqi army, but on anybody carrying a weapon, including (for example) shopkeepers defending their stores against those they perceived as enemy invaders.¹¹⁰ This was arguably in violation of the Geneva Conventions, which rule that civilians may be treated as combatants only in cases where they carry their arms openly and organise themselves into regular armed units of some sort, a standard which is not met by individual civilians protecting their property (Convention III, 1949, article 4 [6]). The Conventions also state that 'In case of doubt whether a person is a civilian, that person shall be considered to be a civilian' (Addition to Protocol I, 1977, article 50 [1]). The attitude of the US government in this matter received condemnation in the press, but there was not the level of public or political outrage required for the USA

¹¹⁰ Synovitz, Ron, *Iraq: Desert Dispatch - Rules Of War Evolving As U.S. Troops Encounter Iraqi Fighters In Civilian Clothes*, www.rferl.org/nca/features/2003/03/28032003152207.asp, 28 March 2003.

to change its policy. This might suggest that, in these enlightened times, an official rule-change concerning the targeting of civilians would not meet with very strong popular resistance, provided the new convention stipulated that only certain sorts of attacks were permitted. For example, a policy permitting attacks on armed and/or hostile civilians, or civilians preparing violent attacks on the enemy (i.e. preparing to become combatants) might not meet with universal disapproval.

It is perhaps impossible to predict with any degree of accuracy how such a change to the conventions of just war might be viewed by the international community, and whether it would be found acceptable to states and their civilian populations.¹¹¹ However, in the absence of good reason to think that we can rule out such a possibility, the question remains of whether this would better limit the inflictions of harms and promote the attainment of goods than adhering to the current convention; and I argue that the answer is no.

Why less limited warfare is not a better moral choice

There are several concerns one might raise about a policy which allowed direct attacks on civilians in certain circumstances. The chief of these is that whatever the ‘certain circumstances’ were, they could be easily abused by warring parties – or at least much more easily than non-combatant immunity, under which states must strive to show that every civilian death was, though regrettable, also militarily necessary. To illustrate this, consider the following report from a journalist in Iraq during the early stages of the recent conflict:

Any Iraqi in civilian clothes that is seen observing U.S. troop movements while carrying a cell phone is considered a legitimate enemy target - an observation post . . . Lieutenant Colonel Charlton said that as the rank-and-file U.S. soldiers become more battle-hardened, they are no longer hesitating to fire at suspicious Iraqis in civilian clothes.¹¹²

¹¹¹ Note that a new convention along the lines indicated above would limit attacks to individual civilians or small groups, and only by ground forces; it would not be reasonable to suppose that a pilot or long-range gunner can discriminate those civilians who are carrying firearms from those who are not and target them accordingly. It would not allow the civilian population to be bombed from the air, for example. Its similarity to existing military practice might allow it to be portrayed as a minor or cosmetic change to the rules.

¹¹² Synovitz, *ibid.*

The American troops reported that many Iraqi soldiers dressed themselves in civilian clothes to avoid being targeted while they prepared to attack the invading troops. Some concealed weapons on their bodies. Therefore civilians who were 'acting suspiciously' were liable to be fired on, whether or not they could be seen to be carrying a weapon. It is true that a new convention allowing attacks on armed civilians could still prohibit this kind of killing, requiring that the civilian was clearly and openly armed and dangerous; but nevertheless, it is reasonable to think that fearful soldiers, who shot first and checked for danger afterwards, would find it very much easier to justify their behaviour than they currently do, and that the practice would proliferate. This possibility of abuse is of especial concern when we consider that there would be no way of punishing states whose armed forces broke the new convention and targeted 'suspected to be armed' civilians. In the absence of any substantive threat of punishment, our best protection against such actions is that killing civilians is so frowned upon.

A second concern, connected closely to the first, is this: if targeting civilians were sometimes justifiable, the horror of war would be multiplied many times, infusing the civilian population with dread and panic. If my country is at war and I know that armed civilians are liable to be targeted by the invading forces, then I will live in constant dread of accidentally appearing to be carrying a weapon of some sort. It will be of little comfort to me that, in theory, I have to be brandishing a recognisable firearm at the enemy to be susceptible to attack. My basic fear will stem from the knowledge that I can no longer be guaranteed immunity by my civilian appearance and by my lack of involvement in the war.¹¹³

It is difficult to anticipate the effects of this kind of widespread civilian fear. A possible good consequence is that states might try harder to avoid going to war, because the toll of warring would become higher. Likewise, politicians (and indeed

¹¹³ For a distressing current example of how fear can be spread through a civilian population by the killing of non-combatants even when this is totally illegal, we can turn once again to Iraq. When the city of Falluja was invaded by US troops in November 2004, NBC reporter Kevin Sites released video footage of US marines shooting an apparently unarmed and wounded Iraqi prisoner in a mosque whom they claimed was 'faking it he's dead'. The sight both chilled and angered many people here in New Zealand; it is hard to imagine the effect on those who viewed it living in Iraq. (See www.kevinsites.net for Kevin Sites' open letter to the Marines with whom he was embedded in Iraq, 21.11.04.)

all civilians) might be more hesitant to campaign for war if it were possible that they would pay the costs themselves. On the other hand, states at war might be more ruthless and bloodthirsty than ever, with the whole population, civilians and all, willing to be fully involved in the war effort. Even those who were previously opposed to the war would, once war was underway, have strong reasons to help their side toward a swift victory, and so states might find themselves with a population 100% behind any military action they cared to perpetrate – an unsettling thought, especially when coupled with the lack of other reasons for restraint. It is difficult to anticipate which of these outcomes would obtain, and to this extent, it is difficult to say whether widespread horror and dread among the civilian population would make war more or less harmful in terms of fatalities and duration.¹¹⁴ However, the horror and dread felt by the civilian population counts as a harm in itself, and we must remember that it would not be ‘compensated for’ by a less fearful class of combatants (who would be as prone to attack as ever). Therefore there is reason to think that allowing some attacks on civilians would, in very simple terms, create more suffering than adhering to non-combatant immunity. This alone is enough to render such a policy *prima facie* objectionable.

Another area in which non-combatant immunity might fare better than some less limited form of warfare is when we consider the overall casualties of war. Recall that I discussed the idea of self-defence in chapters three and four, and concluded that it failed to make a clear distinction between civilians and soldiers. However, it may have a role to play when non-combatant immunity is considered as a utilitarian heuristic, i.e. if it could be shown that killing soldiers is a more effective way of fighting a war than killing civilians, and that fewer lives overall (combatant and non-combatant) will be lost in wars in which civilians are not targeted. This is an empirical question, and evidence for the answer would have to be found by a political scientist of far greater specialty than myself. It is true that soldiers by and large present a threat that civilians do not, but it is also true that governments are reconciled to military and not civilian deaths, and if the pressure on each warring government

¹¹⁴ I must say, however, that the idea that states will be deterred from going to war, simply because their civilians are going to be fearful, sheds a more favourable light on the motivations of states than experience leads one to hold. Financial or territorial interests seem often to outweigh this consideration. If a government’s military is willing and eager to go to war, the opinion of the general public is of discounted importance.

was great enough to bring the war to a swift conclusion, lives might in this way be spared.¹¹⁵ I will leave this argument open, merely pointing out that *if* killing civilians does not help win a war more quickly, then any policy allowing attacks on civilians will generate extra casualties for no extra gain, and would for this reason be undesirable. In the case of ‘armed and dangerous’ civilians, however, it is probable that being allowed to target them would help to win the war more quickly.

How non-combatant immunity has fared

I believe the above arguments give good reasons to think there exists no policy or convention acceptable to states in the near future which will better serve to reduce the infliction of harms, and to promote the attainment of goods, than non-combatant immunity. I conclude that we have a moral obligation to adhere to that convention at this time. However, it does not follow from this conclusion that non-combatant immunity is the best, or even a good, convention for guiding the conduct of states at war. In my following and final chapter, I suggest that other conventions for warfare might one day be acceptable to states, and I explain why and to what extent we have an obligation to agitate for these and for the conditions needed to make them acceptable and functional. In this explanation I concentrate particularly on the need for a system of meaningful international law enforcement. But first, in the remainder of the present chapter, I identify two different tasks for the political philosopher, and discuss the relevance of these tasks to the issues of ‘immediate acceptability’ I have above been discussing.

Section 3: Ideal vs. non-ideal theory

Issues of acceptability

¹¹⁵ Due to the special value we put on civilian lives, it could be that the occasional massive attack on civilians will bring about a swift conclusion to what otherwise would be a lengthy military engagement. As I mentioned above (footnote 92), some have argued that the bombings of Hiroshima and Nagasaki were justified in this way: The US’ show of willingness to inflict massive harm on enemy civilians instantly frightened the Japanese government into surrender, and the bombings caused fewer deaths than the only alternative, a ground invasion of Japan, would inevitably have. The difference is that the deaths caused by the bombings were nearly all civilian, as opposed to the combination of military and civilian deaths which would have been caused by a ground invasion. (The proliferation of nuclear weaponry since 1945 mitigates the application of this argument in a modern setting, as it is no longer the case that a nuclear attack is tolerably likely to meet with surrender. Total war seems the more likely outcome. But it may still be the case that if an attack on civilians is likely to bring the war to a quick end, then such an attack is justified.)

As I have discussed above, the biggest constraint on finding conventions which could be morally valuable *right now* in times of war is their acceptability. Clearly, in terms of avoiding suffering, the ideal convention would require states to use exclusively diplomatic, non-violent and probably multi-lateral means to resolve conflicts. However, as we have seen, those powerful states with economic or political interests in warfare would not be willing to accept this and so the convention could not be made into successful military policy. Even if such a convention were to be adopted on paper, or paid lip service to, violent conflicts would almost certainly erupt when individual states decided they could better their position through force. Therefore such a convention, though at first appealing, would have no real value in solving the moral problems of modern warfare at this point in history.¹¹⁶ It seems, then, that any convention proposed for the conduct of war must stand a good chance of being adopted and followed by the international warring community. This is a general claim about *any* principle or policy being proposed as part of international law, whether it be a convention like non-combatant immunity, which generates a moral obligation but is not based on an intrinsically moral principle, or a proposal which can be defended by appeal to deontological or agent-relative moral considerations.

The above claim, that any convention proposed for the conduct of war must stand a good chance of being adopted and followed by the international warring community, might well sound warning bells in our minds. Do we want to limit our normative claims to such a degree? This question has arisen in other contexts: Kai Nielsen and Allen Buchanan have recently discussed the issue of ideal vs. non-ideal moral theory with reference to the right to secession.¹¹⁷ Below I use their arguments to frame my own.

Must morality be palatable to world leaders?

¹¹⁶ The possibility of meaningful international-law enforcement at this point opens the door to adopting very different types of conventions than those which seem immediately possible. I discuss this more fully both in this and in my final chapter. Also, I state again, there is no reason to suppose that groups of nations could not choose to ratify multilateral treaties more strictly governing the conduct of war, in the absence of a change to international law.

¹¹⁷ Buchanan, A., 'Theories of Secession' (*Philosophy of Public Affairs* 26:1, 1997), pp. 31-61; and Nielsen, K., 'Liberal Nationalism and Secession'. Chapter 6 of Moore, M. (ed.), *National Self-Determination and Secession* (Oxford: OUP, 1998), pp. 127-133.

We should be wary of a statement which seems to say that what counts as a good moral principle must make allowances for what agents are willing to do. Surely what is moral is not constrained by someone's willingness to be moral, but by their capacity. Thus we have *ought implies can*, not *ought implies willing to*.¹¹⁸ The latter would remove any obligation to act against our inclinations. It might seem that saying 'What counts as a good moral principle for the fighting of wars must be agreeable to states' is making this kind of mistake. It might be thought to be placing unacceptable and morally irrelevant constraints on what the moral philosopher is allowed to do. While there is a sense in which this is importantly true, there is another in which it is importantly false. It is true that the moral philosopher ought not to be limited by what people are likely to want to do. It is true that what is right for us to do as individuals does not always correspond to what is palatable. But it is false that ignoring such considerations is appropriate or helpful when formulating the policies of international law which must generate good results when applied, currently without enforcement, across a broad range of cases.

When it comes to formulating laws for the current political climate, with the near-total absence of international law enforcement, it is useless to suggest laws which are not agreeable to powerful states, for the simple reason that they will not be passed or followed. Therefore what is a morally valuable policy *right now* will depend to a great extent on whether states are willing to follow it or not. Suppose that we have to choose between legislating one of three possible ways: policy A, policy B, or no policy at all. Under ideal conditions of total or near-total compliance to the law, A would be the most morally desirable policy (in terms of generating desirable outcomes); but in the actual world the chances of it being made into law or adhered to are very small indeed, and there is no way of enforcing it. Policy B would not be as desirable as A in the ideal world, but stands a good chance of being accepted and followed by the majority of states in the actual one. Legislating no policy at all would be the least morally desirable action in any kind of world, as this licences states to act as they please with no possibility of legal censure from the world arena and no possibility of gradual progress to better policies. It seems clear that, as policy-makers

¹¹⁸ I will assume for the purposes of this discussion that we possess sufficient freedom of will to act morally if we choose; in other words, I will assume a difference between the *can* and the *willing to* of the statement above.

in a non-ideal world, we ought to legislate policy B, while trying to find a way to successfully legislate and enforce A. That B is a concession to the willingness of states to behave a certain way does not render it undesirable, when compared to the viable alternatives.¹¹⁹

This issue of enforcement is one which I will discuss more fully in my final chapter. I will conclude here only that what world leaders are willing to do is, *in the absence of law enforcement*, relevant to which conventions constitute attractive military policy. I agree that this is undesirable, but that is because the absence of enforcement of international law is undesirable. Refusing as a matter of principle to adapt moral edicts regarding policy to the behaviour of states is not a helpful solution to the problem. Although philosophers certainly ought to identify those ideal systems which might one day be implemented, they can do immediate good by forming moral philosophy that is action-guiding in today's world. As I argue below, neither task need preclude the other.¹²⁰

Buchanan on 'minimal realism'

Allen Buchanan, in a 1997 paper on the right to secession and its place in international law, argues for a criterion of 'minimal realism' in all our proposals for international law. A theory counts as minimally realistic just in case 'it has some significant prospect of eventually being implemented through the actual processes by which international law is made and applied.'¹²¹ The reason our policies must be minimally realistic is, says Buchanan, because this is the only way they will do any good. This is not to sacrifice morality for acceptability; rather, Buchanan argues that reforming international law is best achieved by introducing principles which are in accord with its already deeply entrenched and morally progressive beliefs:

¹¹⁹ This is an echo to my statement in chapter 3 that the utilitarian ought not to be blamed for permitting seemingly-bad actions in circumstances in which the background misapplication of the moral theory means that there are no preferable actions available. See footnote 65.

¹²⁰ We must remember that the comparison we are making, when choosing to whether to legislate a certain convention (agreeable to states) into international law, is between the convention and the way things currently are, not between that convention and the way things might ideally be. If adopting the convention will improve on the status quo, then that is a *prima facie* reason to follow it, no matter how it compares to the ideal situation. (As I discuss below, it will often be the case that we ought also to work to have a better convention adopted as law, but the obligation to try to do something even better in the future does not remove the obligation to do something good in the meantime.)

¹²¹ Buchanan, p.42.

Moral theorising... can provide significant guidance for international legal reform only if it coheres with and builds upon the most morally defensible elements of existing law, but...noninstitutional moral theories [those which require fundamental changes to the structure of international law] fail to satisfy this condition.¹²²

Buchanan concludes that any moral policy which would require fundamental changes to the way our system of international law operates is so unlikely to meet with success that its very proposition is an idle exercise. If we want to do good in the real world, we ought to identify those policies which conform to the basic structure and ideals of international law, but are 'morally progressive' in that they propose some small and achievable change for the better. (I further consider these arguments in my final chapter, where I suggest that the appeal of his standard of minimal realism rather depends on our interpretation of his use of the word 'eventually'.)

Must morality seek to accommodate the political status quo?

Kai Nielsen, in a 1998 paper replying to Buchanan, rejects the *realpolitik* he takes Buchanan to be espousing. Nielsen argues that rather than accommodating the lawmaking processes of the status quo, moral philosophers ought to be

articulating an ideal theory, for liberal democratic societies, which involves thinking institutionally about what international law...would look like under ideal conditions, i.e. under conditions in which states in such societies actually behaved in accordance with the moral principles embedded in the very ideals of liberal democratic society.¹²³

That is, rather than accepting that we must tailor our recommendations for international law to the realities of today's political setup and the whims of powerful states, we ought to be identifying what that law should ideally enforce, according to its own moral foundation in liberal democratic ideals: 'The moral tail should wag the legal dog'.¹²⁴

¹²² Buchanan, p. 32

¹²³ Nielsen, p. 130

¹²⁴ Nielsen, p. 129

However, I suggest that Nielsen's claim – that the current political setup of the world should be irrelevant to what is morally desirable – is only true according to an illegitimately limited understanding of the proper aims of moral philosophy. It is true that if we are 'articulating ideal theory', or trying to identify the *best possible* political system under which human beings might live, we ought not to be influenced by the fact that the USA is currently the strongest world power, or that communist states have fared badly in recent history. However, if we are deciding whether something ought to be part of international law as it is currently structured and implemented, then facts about the current political world are highly relevant. If we ignore them, then we will end up with recommendations for international law which require impossibly quick fundamental changes to world political systems. It is quixotic and unreasonable to think that any such recommendations could create a desirable outcome.¹²⁵ If they were passed and followed, they may well be followed to ill effect; but they are most likely simply to be rejected or ignored.

I strongly believe that determining whether or not a policy ought to be made into international law as it is currently structured is a legitimate project within moral philosophy. Although Nielsen rightly identifies the need to have an ideal toward which to aim, it must be seen as a shortcoming of his arguments that they prevent us from advising the lawmakers of today about what the law ought to be.

Two projects for political philosophy

Buchanan and Nielsen each argue compellingly for the value of two distinct tasks within moral philosophy. However, I believe that the theoretical and practical success of these two projects is codependent, and thus that it is wrong to take one of them to be important at the exclusion of the other. One of my aims in these last chapters is to show that there is moral-theoretical space, at least within a utilitarian framework, for both types of task to create real moral obligations.

The first task, which I think Buchanan dismisses too quickly but which Nielsen perhaps overemphasises, is to identify and argue for the *best* kind of system under

¹²⁵ Buchanan argues convincingly for this, p. 44ff.

which human beings can live. This kind of theorising ought to take into account unchanging facts about human beings, but arguably not contingent and arbitrary facts of human history (except as they provide information about human psychology). This is because when we, as moral theorists, try to imagine what political organisation will look like in the best possible world (whether it be the ideal utilitarian world or the ideal Kantian world or the ideal virtue-ethical world), our thinking ought to be constrained only by those features of human existence which we know must be present in such a world, i.e. our physical and cognitive limitations. The value of this project lies in giving us a goal to aim toward, and in identifying which kinds of decisions, principles and systems help us work toward this 'ideal society' and which do not. (When Buchanan requires that policies be 'morally progressive', he must, after all, think there is some way to determine toward what they ought to progress.) From a utilitarian perspective, we must identify those ideal political systems which best maximise utility within the constraints of human existence in order to safeguard the long-term utility of our actions. If our view is always limited to the short-term generation of utility, we risk forgoing massive and enduring long-term benefits for more minor immediate ones.

The second project of moral philosophy, which Nielsen seems to consider misguided, is to identify and argue for ethical policies in a way which is action-guiding for those currently in power; and here I accord with Buchanan in thinking that large concessions to the status quo are essential. This project will necessarily take into account the nature of the current political world, as well as its sociocultural and economic history. Its value lies in finding ways for the ideals identified by the first project to be eventually be reached; and so the value of each project is connected to the value of the other. It is useless to have an ideal with no means of reaching it, but it is likewise useless to try to make the world better without an idea of what that would mean.¹²⁶

¹²⁶ Nielsen at one stage seems to acknowledge this in his mention of John Dewey's 'means-end continuum'. He states that 'we would not only have to consider the ideal, but the probabilities and conditions for attaining or at least approximating what the ideal calls for and the costs of such an attainment' (p. 131). But to do this is, I think, to do what he argues we should not do: develop normative recommendations which are influenced by the status quo. It might be that he considers the 'means-end' moral task as subsumed under the ideal one; but still, as long as there are better and worse ways of working toward an ideal, we need some way to tell the better from the worse, and Nielsen seems unable to give guidance here.

As an illustration of the difference between the projects, and the way in which they are interrelated, we might once more consider the pacifist position, this time with reference to arguments commonly made against it. One criticism often levelled toward pacifists is that they are unrealistic about the possibility of peace and therefore their philosophy is useless.¹²⁷ Wars, their critics say, are obviously going to be fought in the foreseeable future. If the pacifist refuses to tell us how we might best fight them, simply labelling them all as immoral, then isn't she missing the chance to do some real-world good? Surely she would best promote the good by finding ways to make inevitable wars less evil.

I think the pacifist can agree with this last statement without relinquishing her commitment to pacifism. She might say something like this: 'We best promote the good by acknowledging that all war is wrong, as shown by pacifist arguments a, b and c; and by acknowledging that war actually can be avoided, as shown by x, y and z. In the meantime, however, we ought to minimise the infliction of harms generated by war by acting in accordance with principles p, q and r.' The goal of trying to limit the harm caused by up-and-coming wars need not work against the goal of putting an end to all war.

With reference to non-combatant immunity, and to clarify the project I have undertaken here, it is important to realise that the statement 'Wars in which civilians are not targeted are preferable to wars in which they are' does not entail the truth of the statement 'A war in which civilians are not killed is a just war'. Likewise, 'action x is morally preferable to action y' does not automatically entail that action x is the right thing to do, so long as there are other options to be considered. The pacifist argues that there *is* another option to be considered when we talk about war: Rather than wars in which only combatants are killed, or wars in which anyone can be targeted, it is possible to have no war at all, and this alone is morally preferable in comparison to every alternative and thus deserves the judgement 'morally right'.¹²⁸

¹²⁷ See Cheyney Ryan's excellent article 'Self-Defense, Pacifism, and the Possibility of Killing' (*Ethics* 93, 1983, pp. 508-524) for a delightful survey of pacifism-bashing quotes from contemporary moral philosophers (p. 509).

¹²⁸ The success of the pacifist argument relies on the claim that war is not a necessary part of human existence. If there is no possible political system under which wars will not take place, then the pacifist

I will not discuss the tenability of the pacifist position here. It is enough to recognise we get no guarantee of enduring moral rightness by adhering to that standard which is preferable to its currently feasible alternatives, given that this feasibility is very much limited by contingent political facts about the world. In order for adherence to non-combatant immunity to be a lasting guarantee of moral rightness, there would have to be no prospect of future preferable alternatives. In other words, we would have to be (at least with respect to the conduct of war) living in the best possible political society, the society than which none could generate a greater balance of happiness over suffering. I see no reason to think that this is the case; on the contrary, there is good reason to fervently hope that it is not.

A revision of the Just War theorist's aims

We can conclude from the above discussion, and from Nielsen's criticisms of Buchanan, that there is value in identifying conventions, principles and policies which might be part of international law in the future even if they do not seem immediately acceptable to powerful states. In the following, final chapter I argue that for the further claim that we have a moral obligation to identify how to best reach these moral goals.

My conclusion as regards non-combatant immunity is, as I discussed above, that states have a real moral obligation to adhere to the principle of non-combatant immunity, despite the fact that it does not point to any intrinsic moral difference between civilians and soldiers. This obligation stems from the fact that wars fought according to non-combatant immunity are less harmful than they would be if states did not adhere to the principle, or if they adhered to any other immediately acceptable principle governing the conduct of war. However, as I discuss in the next chapter, there are good reasons for thinking that even better moral conventions could be adopted if we made appropriate changes to international law and its enforcement. As long as these better conventions are *possible* conventions, we have an obligation to

goal really is unattainable. A thorough evaluation of the plausibility of various pacifist claims is outside the scope of this thesis, although (as I explain in chapter 6) I see no reason to think their general goal is unattainable. It is worth mentioning, though, that the waging of *economic* war, fuelled by the uneven spread of natural resources in the world, strikes me as of equal concern to the pacifist as traditional military warfare, insofar as it contributes to massive innocent suffering.

motivate for their adoption. Our concern ought to be in creating a situation in which 'ideal' morality can be enforced, not in rejecting it wholesale for something more immediately feasible.

What this ideal morality will be depends on how we judge the evils of war, and our rejection of non-combatant immunity importantly influences this. When we discard the idea that a civilian death is to be more stringently avoided than a soldierly one (or, to phrase the comparison the other way, when we realise that that a soldierly death is no more acceptable than a civilian one), we find that the harm to be avoided in war relates to the *overall* carnage and loss of life, rather than the carnage and loss of life within a particular group. Therefore when we look at conventions which might guide the conduct of war, rather than endorsing something like non-combatant immunity, we are more likely to search for conventions which allow wars to be fought with a view to limiting *all* death and suffering.

CHAPTER SIX

Concluding remarks

Identifying our obligations

The central purpose of this thesis has been to explore and identify our moral obligation with regard to acts of war, and in particular with regard to the standard of non-combatant immunity, a traditional part of *jus in bello*. I have argued that while we do have an obligation to adhere to this standard, the obligation is not generated by any intrinsic moral properties of soldiers and civilians. Rather it is generated by the lack of preferable alternatives in the short term, and relies on utilitarian moral reasoning. A rejection of utilitarianism thus implies a rejection of the only moral-theoretical basis for non-combatant immunity. I conclude that immunity theorists must either be utilitarian, or they must relinquish the claim that their principle is morally defensible.

As I hinted in chapter 5, that non-combatant immunity can only be defended by utilitarian reasoning has important implications for the future of the principle as an arbiter of rightness. To adhere to non-combatant immunity today is right because no preferable alternative exists, but it does not follow that to adhere to non-combatant immunity forevermore will ensure that all our future actions are right. A utilitarian conception of *jus in bello* will require that its principles change with changing political possibility, and what is politically possible in the future must rely heavily on what we do and think and say today. Because the utilitarian standard of rightness requires us to consistently aim for the *best possible* action in the circumstances, then if we can act today so that innocent soldiers and innocent civilians need not be killed in war in the future, we have an obligation to do so. In fact, we have an obligation to work toward implementation of the best possible policy regarding the conduct of war, and as I argue further below, it is difficult to imagine what the limits on that might be.

There are, I think, three important questions which I have not considered here, and which must be answered before the nature and extent of our current and future

obligations concerning *jus in bello* in general, and non-combatant immunity in particular, can be understood. These three questions are:

- 1) What are the potential conventions, principles or policies governing the conduct of war and political relationships?
- 2) How do we judge the desirability of possible policies?
- 3) What political and social changes need to occur in order that moral progress is made toward the best possible policy or policies?

I consider each of these questions briefly below, not in any attempt to answer them, but rather in an attempt to clarify the questions and indicate where further work might usefully be done.

1) What are the potential conventions, principles or policies governing the conduct of war and political relationships?

I stated in chapter 5 that the political philosopher has two tasks: to decide what states ought to do in the near future, and to decide how states would behave in an ideal world. The former task is heavily influenced by facts about the current political climate, but the latter is bounded only by human possibility, and history teaches us that this excludes very little. This is why I stated that Buchanan's requirement that policies have a chance of being 'eventually' enforced is so problematic; interpreted liberally, it doesn't seem to be a restriction at all. In chapter 2 I made brief mention of the historical receptions accorded suffragettes, anti-nuclear demonstrators, and black civil rights campaigners. I invoke them again here because they are illustrations of the startling difference in moral beliefs between today's society and the societies of 2000, 500, 50 or even 20 years ago, and they tell us something interesting and important about our own 'traditional' moral views. If in 1850 you had told an early settler of New Zealand that, in 2001, our Prime Minister, Leader of the Opposition, Chief Justice and Governor General would all be women, his response would likely be outright incredulity. Likewise, if someone had claimed in 1994 that, ten years later, a central city supermarket would bow to public pressure and cease to stock battery-farmed eggs, they would have sounded like an idealistic dreamer; and yet they would have been right. Our best evidence tells us we are but one stage in a long march of

moral change; and given that there are features of our everyday moral thinking which would strike our ancestors as outrageously quixotic (or perhaps outrageously immoral), we would be arrogant to assume that our prevailing moral beliefs are here to stay.

What this suggests it that we may not have a very accurate idea of the moral beliefs, and corresponding political possibilities, of the world a hundred years from now. This should, I think, lead us to be liberal in our demands, rather than conservative, because in assuming that great change is possible, we thereby make it more likely. The suffragettes, anti-animal-cruelty campaigners, and black civil rights activists of the past would have achieved little if they had taken what would have struck most people of the time as a realistic view of their chances of attaining their objects. In the incessantly quoted words of Margaret Mead, 'Never doubt that a small committed group of individuals can change the world. Indeed it is the only thing that ever has.'¹²⁹ There is much to be said for the visionary, so long as there are discernible steps which may be taken toward the vision.

Therefore when we come to question what sorts of policies might one day be successful in governing war, it is not ridiculous to think of true diplomatic solutions, or a meaningful world court of justice, or practical global economic interdependence, or a real ethical commitment on the part of the world's population to peaceful conflict resolution. It is not unreasonable to suppose that war could one day be made unprofitable and thus undesirable even to those without ethical qualms about warring. Robert Holmes points out that we commit billions of dollars and millions of lives every year to the war industry; he asks us to imagine what might be possible if we were willing to expend these resources in non-violent conflict resolution instead.¹³⁰

The shifts in attitudes and political organisation suggested here may not be any more profound than the kinds of shifts the world has already seen over and over again. The most important thing in considering future policy regarding war is not to be keenly

¹²⁹ Mead, M., 'Coming of Age in Samoa' (New York: New American Library, 1949), quoted at www.transculturalcare.net/Sayings.htm.

¹³⁰ Holmes' discussion of the economics and politics of the 'war system' can be found on p. 268ff of *On War and Morality*. Holmes accuses his detractors of having a 'lack of imagination about the potential of nonviolence', p. 270.

conscious of our limitations, but rather to be certain about the desirability of the direction of change. For this, we need accurate information about what is bad about war, which leads me to my second question:

2) How do we judge the desirability of possible policies?

‘War is hell’, writes Michael Walzer, and proceeds to examine just what kind of hell war is, and how we can diminish its harms.¹³¹ However, Walzer’s commitment to non-combatant immunity places a limit on his project which I have argued there is no moral reason to place. If we accept that there are no necessary intrinsic moral differences between civilians and soldiers, then we should accept that (contingent facts about societal attitudes aside) there is no reason to differentiate between the badness of civilian deaths and the badness of soldierly deaths when we judge the evils of war. Rather, we ought to aim to prevent wars when possible, and when not possible, to minimise the suffering of combatants and non-combatants alike.

This new way of thinking about the harms of war signifies a need for a major change to *jus in bello*. Rather than the standard of acceptable killing in war being identified by non-combatant immunity, ‘military necessity’ and concerns about proportionality take on new significance. If an act of war can be reasonably expected to save more lives than it squanders, and protect more property than it destroys, and there is no better route to the same desirable conclusion, then it here is good reason to think that it may be justified, *even if it involves the direct targeting of civilians*. Thus if it can be shown that by not respecting non-combatant immunity, a state can fight a war more effectively and thus limit overall carnage, there do not seem to be good moral reasons for preventing them from doing so. This conclusion, however, relies on a number of preconditions which do not currently hold, particularly that there is an international law-keeping body capable of preventing a country from fighting unjustly, or punishing it when it does. Also, it must be remembered that, as I discussed in chapters 4 and 5, targeting civilians will generally not be as effective in bringing a swift conclusion to war as targeting soldiers. In addition, due to the comparatively high social cost of civilian deaths, any good utilitarian will agree that if a society is more

¹³¹ Walzer, M., p. 24ff.

willing to sacrifice their soldiers, and the soldiers more willing to be targets, then targeting the soldiers rather than the civilians may be preferable in the immediate future. However, this characterisation of our values does highlight the callous way non-combatant immunity can lead us to view soldiers: as expendable weapons of war, rather than moral agents possibly as deserving of protection from attack as we are ourselves.

3) What political and social changes need to occur in order that moral progress is made toward the best possible policy or policies?

I stated above that it is difficult to imagine the nature and degree of moral and social change that might occur in the century and beyond, and so it might seem that to try to take steps toward a better world is a futile enterprise. If we cannot know what change is possible, then we cannot know what goals will be achievable and when. However, I think there is a lot we do know: we know why states go to war, and we know what deters them from warring. If we can eliminate the former conditions whilst enforcing the latter, we stand a good chance of making moral progress.

So, very briefly: states go to war to make money, or to assert or defend their culture, religion or political beliefs, or both. They are deterred from warring when war is too expensive, whether that be in financial, political or social terms. War can be made financially undesirable by encouraging many economies to be inter-reliant; by controlling the arms trade such that weaponry is costly; by de-glamourising the righteousness of soldiery so that soldiers require high pay as an inducement to fight; by placing economic sanctions on states who are waging unjust wars; and so on. War can be made politically or socially undesirable by educating the global population to be peace-loving and tolerant of difference; by empowering minority groups so that their voices are heard; by granting the means for political groups to vie for power in a fair and controlled manner; and by instituting an effective means of intervening and ending unjust wars, punishing aggressive states, and compensating the victims of their aggression. It may also be made undesirable by encouraging the spread of active

democracy: in 1795, Kant argued in 'Perpetual Peace'¹³² that democracy favours peace, and this claim has been endorsed by Margaret Thatcher, Bill Clinton, George Bush Jnr and many modern pacifists.¹³³ Kant's argument is that in a democracy, 'the consent of the citizens as members of the State is required to determine at any time the question, 'Whether there shall be war or not?''', and peace is thus favoured because 'nothing is more natural than that [the citizens] should be very loth to enter upon so undesirable an undertaking; for in decreeing it they would necessarily be resolving to bring upon themselves all the horrors of War.'¹³⁴ In addition, we can expect citizens committed to protecting peace to elect governments who assert a similar commitment.

I do not think that any of the above means of preventing war are impossible. However, many of them rely on a kind of multilateral international law-enforcement which we do not have today, one strongly focussed on punishing the violent and protecting the peaceful, and capable of balancing the interests of states vastly dissimilar in power and influence. The success of such a system will rely on the co-operation of many different governments, and those policies which will be acceptable to governments in the future depend largely on what pressures are brought to bear on those currently in power. This in turn is influenced by what the general population is motivated or encouraged to ask for and consider valuable; and thus the first step toward multilateral international law-enforcement aimed at preventing war and protecting peace is to recognise and teach that war is a serious harm, and that there are alternatives to war which need not result in the triumph of the powerful.

¹³² Kant, I., 'Perpetual Peace: A Philosophical Sketch', 1795, trans. W. Hastie. In *Kant's Principles of Politics, including his essay on Perpetual Peace. A Contribution to Political Science* (Edinburgh: Clark, 1891), pp. 77- 149.

¹³³ Margaret Thatcher, on a visit to Czechoslovakia in 1990, famously stated that 'democracies don't go to war with one another' (quoted in *The Economist* April 1995, p. 18), and a study published in *Peace* magazine in 1999 by Rudolph J. Rummel seems to support this claim, at least amongst modern democracies. Rummel showed that of the 353 pairs of nations engaged in 'war' (as precisely defined according to various criteria) between 1816 and 1991, no war was between two 'democracies' (likewise precisely defined). While the rigorous nature of these definitions mitigates his absolute conclusion somewhat, nevertheless Rummel's findings are striking. An excellent summary of his study can be found in a BBC news story of November 17, 2004, at http://news.bbc.co.uk/2/hi/uk_news/magazine/4017305.stm, and the original study at www.peacemagazine.org/archive/v15n3p10.htm.

¹³⁴ Kant, pp. 89-90.

I hope to have shown in this thesis that part of what it means to recognise that war is a serious harm is to recognise that even the deaths of combatants are to be regretted, in many cases as fervently as are the deaths of civilians. It is to recognise that even wars in which no civilians are killed may represent the massive infliction of unjust harm, and that strict adherence to *jus in bello* does not provide states with eternal protection from wrongdoing. Given the history of moral progress I discussed above, I see no reason to consider it impossible that there ever be this kind of change in popular thinking about war, and I believe that such a change represents our best hope of encouraging political dedication to peace.

Bibliography

Books and articles:

1. Buchanan, A., 'Theories of Secession'. *Philosophy of Public Affairs* vol. 26 no. 1, 1997, pp. 31-61.
2. Crawford, N. C., 'The Slippery Slope to Preventive War' (*Ethics and International Affairs* 'Roundtable on Evaluating the Preemptive Use of Force', vol. 17, no. 1, 2003).
3. Herman, B., *The Practice of Moral Judgement* (Cambridge, Ma: Harvard University Press, 1993).
4. Holmes, R., *On War and Morality* (Princeton, NJ: Princeton University Press, 1989).
5. Lichtenberg, J., 'War, Innocence, and the Doctrine of Double Effect'. *Philosophical Studies* vol. 74, 1994, pp. 347-368.
6. Mavrodes, G., 'Conventions and the Morality of War'. *Philosophy and Public Affairs*, vol. 4 no. 2, 1975, pp. 117-131.
7. Mead, M., *Coming of Age in Samoa* (New York: New American Library, 1949).
8. Nielsen, K., 'Liberal Nationalism and Secession'. Chapter 6 of Moore, M. (ed.), *National Self-Determination and Secession* (Oxford: OUP, 1998), pp. 127-133.
9. Norman, R., *Ethics, Killing and War* (Cambridge: CUP, 1995).
10. Nozick, R., *Anarchy, State and Utopia* (Oxford: Blackwell, 1974).

11. Otsuka, M., 'Killing the Innocent in Self-Defence.' *Philosophy and Public Affairs*, vol. 23, no. 1, 1994, pp. 74-94.
12. Primoratz, I., 'Michael Walzer's Just War Theory: Some Issues of Responsibility'. In *Ethical Theory and Moral Practice*, vol. 5, no. 2, June 2002, pp. 221-243.
13. Rodin, D., *War and Self-Defense* (Oxford: OUP, 2002).
14. Rummel, R. J., 'Democracies Don't Fight Democracies'. In online magazine *Peace Magazine*, May-June 1999. Retrieved from www.peacemagazine.org/archive/v15n3p10.htm.
15. Ryan, C., 'Self-Defense, Pacifism, and the Possibility of Killing'. *Ethics*, vol. 93, 1983, pp. 508-524.
16. Smart, J. J. C., and Williams, B., *Utilitarianism: For and Against* (Cambridge University Press, Cambridge, 1973).
17. Synovitz, R., *Iraq: Desert Dispatch - Rules Of War Evolving As U.S. Troops Encounter Iraqi Fighters In Civilian Clothes*, 28 March 2003. Retrieved from www.rferl.org/features/2003/03/28032003152207.asp.
18. Thompson, J. J., 'Self-Defense and Rights', *The Lindley Lecture* (Lawrence, Kansas: Lawrence University of Kansas Publications, 1967).
19. Uniacke, S., *Permissible Killing: The Self-Defence Justification for Homicide* (Cambridge: CUP, 1994).
20. Walzer, M., *Just and Unjust Wars 2nd ed* (New York: Basic Books, 1992).

Online news stories and resources:

21. A reference guide to the Geneva Conventions published online by the Society of Professional Journalists, 2002: www.genevaconventions.org.
22. CBS News story about desertions of American soldiers, December 8, 2004: www.cbsnews.com/stories/2004/12/06/60II/main659336.shtml.
23. Information about the plight of child soldiers all over the globe: www.child-soldiers.org.
24. Information about the surprising existence of Flat Earth societies: www.flat-earth.org/society/about.html, www.lhup.edu/~dsimanek/fe-scidi.htm, and www.alaska.net/~clund/e_djublonskopf/Flatearthsociety.htm.
25. Kevin Sites' open letter to the Marines with whom he was embedded in Iraq, November 21, 2004: www.kevinsites.net.
26. *Los Angeles Times* story about the US public's attitudes to the war on Iraq, December 17, 2002: www.latimes.com/news/custom/timespoll/la-na-iraqpoll17dec17.story.
27. *Newsmax* story on the Ken Bigley hostage crisis, December 20, 2004: www.newsmax.com/archives/articles/2004/9/22/162522.shtml
28. *Philadelphia Inquirer* story about the US public's attitudes to the war on Iraq, June 14, 2003: www.philly.com/mld/inquirer/news/front/6085261.htm?1c.
29. PIPA (Programme on International Policy Attitudes) story about the US public's attitudes to the war on Iraq, June 2003: www.pipa.org/whatsnew/html/new_6_04_03.html.
30. Andrew McCaw's *Scrotology*, a theory of everything: forthcoming.

31. Times Online story about desertions of American soldiers, December 10, 2004: www.timesonline.co.uk/printFriendly/0,,1-3-1397131-3,00.html.
32. BBC News story summarising Rudolph J. Rummel's study on democracy and peace, November 17, 2004: http://news.bbc.co.uk/2/hi/uk_news/magazine/4017305.stm. The original study: www.peacemagazine.org/archive/v15n3p10.htm.