

8. Proportionality

The law of armed conflict prohibits attacks “which may be expected to cause incidental loss of civilian life, injury to civilians, damage to civilian objects, or a combination thereof, which would be excessive in relation to the concrete and direct military advantage anticipated.”¹ This legal norm reflects the moral principle of *jus in bello* proportionality, and is generally considered (along with noncombatant immunity) one of the two fundamental principles of the law governing the conduct of hostilities. The effect of this norm is to make lawful acts that the agent knows will, or foresees may, inflict death or serious physical injury on human beings who have not forfeited their relevant moral or legal rights. Since such acts infringe fundamental moral rights, any legal norm permitting such acts bears a heavy burden of moral justification.

Moreover, an account of *jus in bello* proportionality must satisfy two apparently conflicting demands. Such an account must explain how we can rationally compare civilian losses with military advantages. At the same time, the law applies to all attacking forces *symmetrically* and *independently* of the *jus ad bellum* morality or legality of their overall war effort. Existing accounts of *jus in bello* proportionality satisfy either one demand or the other. In this chapter, I offer a new account that satisfies both demands. Along the way I hope to answer a number of additional questions that any serious account of *jus in bello* proportionality must address.

Incommensurable Values

On its face, *jus in bello* proportionality seems to call for a comparison between *incommensurable* values, that is, values irreducible to any common measure by reference to which their instantiations could be ranked.² To be sure, we often make confident proportionality judgments in extreme cases. For example, the ICRC writes that ‘the presence of a soldier on leave obviously cannot justify the destruction of a village. Conversely, if the destruction of a bridge is of paramount importance for the occupation or non-occupation of a strategic zone, it is understood that some houses may be hit,

¹ Protocol Additional to the Geneva Conventions of 12 August 1949, and Relating to the Protection of Victims of International Armed Conflicts (Protocol I) art. 51(5), June 8, 1977, 1125 U.N.T.S. 3 (entered into force Dec. 7, 1978).

² See, e.g., [MAJ. GEN.] A.P.V. ROGERS, LAW ON THE BATTLEFIELD 17 (1996) (‘Some delegations at the diplomatic conference at which Protocol I was negotiated ... were reluctant to include any reference to the proportionality rule because of the difficulty of comparing things that were not comparable (i.e. military advantage and civilian losses)’); Michael N. Schmitt, *The Principle of Discrimination in 21st Century Warfare*, 2 YALE HUM. RTS. & DEV. L.J. 143, 151 (1999) (‘Optimally, balancing tests compare like values. However, proportionality calculations are heterogeneous, because dissimilar value genres—military and humanitarian—are being weighed against each other. How, for example, does one objectively calculate the relative weight of an aircraft, tank, ship, or vantage point in terms of human casualties?’).

but not that a whole urban area be levelled [sic].³ Similarly, the Israeli High Court writes that ‘shooting at [a sniper firing on soldiers or civilians] is proportional even if as a result, an innocent civilian neighbor or passerby is harmed. That is not the case if the building is bombed from the air and scores of its residents and passersby are harmed’.⁴

Our confident judgments in such extreme examples suggest that even if military advantages and civilian losses are *incommensurable* they are not fully *incomparable*.⁵ At the same time, our less confident judgments in close and intermediate cases suggest that military advantages and civilian losses are at most imprecisely or roughly comparable. On this view, we should not expect *jus in bello* proportionality to provide clear or enforceable guidance in the vast majority of tactical situations that attacking forces confront.

Henry Shue proposes that we can improve our proportionality judgments by sorting particular military advantages and civilian losses into rough categories along the following lines:

	Military Advantage	Civilian Losses
Level 1	Important	Moderate
Level 2	Compelling	Severe
Level 3	Decisive	Tragic

According to Shue, ‘it is excessive to inflict civilian losses of a category higher than the category of military advantage anticipated’.⁶ On this view, an attacking force may inflict moderate civilian losses in pursuit of an important, compelling, or decisive military advantage; severe civilian losses in pursuit of a compelling or decisive advantage; and tragic civilian losses only in pursuit of a decisive military advantage.

Shue does not attempt to fix the boundaries of each category, but he imagines that the task of doing so would proceed along parallel tracks, with military experts categorizing military advantages according to military standards and individuals with ‘humane sensibilities’ categorizing civilian losses according to moral standards. Independent moral judgment would then be exercised to sort the categories created by the two groups into three

³ Commentary on the Additional Protocols of 8 June 1977 to the Geneva Conventions of 12 August 1949, para. 2214 (Yves Sandoz et al. eds., 1987) [hereinafter Protocol I Commentary].

⁴ HCJ 769/02, Pub. Comm. Against Torture in Israel v. Gov’t of Israel, Judgment, Dec. 11, 2005, para. 46, available at http://elyon1.court.gov.il/Files_ENG/02/690/007/A34/02007690.A34.pdf. See also HCJ 2056/04, Beit Sourik Village Council v. Israel, Judgment, June 30, 2004, para. 41, available at http://elyon1.court.gov.il/files_eng/04/560/020/a28/04020560.a28.pdf.

⁵ See Ruth Chang, *Introduction*, in *INCOMMENSURABILITY, INCOMPARABILITY, AND PRACTICAL REASON*, 14 (Ruth Chang ed. 1997)(arguing that incommensurable values sometimes permit ‘nominal-notable’ comparisons in extreme cases).

⁶ Henry Shue, *Indiscriminate Disproportionality*, manuscript at p.24.

(or more) levels. Each level would then contain categories of losses and advantages that are roughly equal in moral weight.

The virtue of Shue's approach is that it can apply symmetrically to all sides of a conflict independently of the justice of their respective war aims. The basic shortcoming of Shue's approach is that it does not add anything to our intuitive judgments. Since it does not identify the moral principles and empirical assumptions underlying our judgments, it cannot tell us when those principles are inapplicable or those assumptions are unsound. Since it does not explain the basis of our intuitive proportionality judgments in extreme cases, it cannot help us make inferential proportionality judgments in non-extreme cases. This approach cannot tell us what to consider or ignore when we make such judgments, or what circumstances might render our intuitive judgments more or less reliable. In short, Shue's approach does nothing to illuminate our intuitive judgments but simply leaves them as he found them.

On a practical level, Shue's approach organizes our intuitive moral judgments but cannot replace or improve upon them. Certainly, we can rank civilian losses by their moral gravity and military advantages by their military utility, but we lack a common standard of value through which the two rankings can be integrated. True, we could directly judge one item in the first ranking comparable to one item in the second ranking, and use this direct comparison as an anchor for integrating the remainder of the two rankings. However, simply ranking two items does not reveal the degree of difference between them. It is not enough to know that a "compelling" military advantage is greater than an "important" military advantage and that "severe" civilian losses are greater than "moderate" civilian losses. We also need to know *how much* greater advantages or losses in one category are than advantages or losses in another category. Only then can we use our anchor point to generate parallel rankings of imprecisely comparable values. It seems unlikely that Shue's approach could generate such a complete ranking.

Military Advantage and Just Cause

Thomas Hurka and Jeff McMahan reject the view that civilian losses and military advantages are incommensurable values on the grounds that military advantages have no *intrinsic* value at all. Indeed, soldiers ought not kill opposing combatants, capture strategic territory, or destroy military equipment for its own sake. Moreover, such military advantages have *instrumental* or *derivative* value only if they contribute to some further, intrinsically valuable state of affairs. Importantly, the defeat of an opposing armed force has intrinsic value only if one fights for a just cause, that is, a war aim (such as national self-defense or humanitarian intervention) that morally justifies resorting to or continuing the use of military force.

Hurka and McMahan conclude that the value of a military advantage, if any, lies in the contribution it makes to the achievement of a just cause. Conversely, a military advantage that contributes to an unjust cause has no moral value. Hurka and McMahan therefore reject the independence of *jus ad bellum* just cause and *jus in bello* proportionality, and with it the symmetrical application of *jus in bello* proportionality to just and unjust combatants alike.⁷ On their view, unjust combatants cannot conform to *jus in bello* proportionality. A military advantage cannot inherit moral value from a war aim that has no moral value to pass on, and a military advantage with no moral value cannot morally justify inflicting civilian losses. On their view, every harm unjust combatants inflict on civilians is morally disproportionate.

For just combatants, this moral standard makes for an impossible decision procedure. Soldiers would first have to measure the moral importance of their war aims, since a similar contribution to a more important war aim would justify more extensive civilian losses.⁸ Soldiers would then have to measure the degree to which the achievement of a particular military advantage would increase the probability of achieving their war aims, discount the value of their war aims by this marginal increase in probability, and compare the resulting expected value with the civilian losses they expect to inflict. Finally, if their just cause depends on collective values such as national self-determination, soldiers will have to somehow compare incommensurable (individual and collective) values. Nor is it clear how a more manageable decision procedure could be derived from the moral standard proposed.

Moreover, on this view, proportionality prohibits but cannot regulate the conduct of unjust combatants. This view provides no moral guidance to combatants who are forced to fight without a just cause. To be sure, unjust combatants can still choose to minimize the harm they inflict on civilians in pursuit of their war aims, and even to place themselves at greater risk to reduce the risks they impose on civilians. But since unjust combatants cannot pursue their unjust war aims by *proportionate* means it is useless for them to try. Nor can observers condemn a particular military operation simply by comparing the military advantage it achieves with the civilian losses it inflicts. One cannot judge particular operations on their own terms; one can only protest the war as a whole.

Hurka and McMahan offer an account of *jus in bello* proportionality that is morally intelligible but applies asymmetrically and yields an

⁷ Put another way, *jus in bello* proportionality is perfectly continuous with *jus ad bellum* proportionality. Just as *jus ad bellum* proportionality justifies civilian losses inflicted by the war as a whole by reference to the importance of a just cause, *jus in bello* proportionality justifies civilian losses inflicted by a particular military operation by reference to the contribution of that operation to the achievement of the same just cause.

⁸ See Thomas Hurka, *Proportionality in the Morality of War*, 33 *PHILOSOPHY & PUBLIC AFFAIRS* 34 (2005); JEFF MCMAHAN, *KILLING IN WAR* (2009).

impractical decision procedure. Shue offers an account that applies symmetrically but is morally inexplicable and of little practical use. In the following sections I will try to do better.

Inflicting and Preventing Harm

As we saw in Chapter 2, killing and injuring others presumptively infringes their basic moral rights. In some cases, these others may, by their voluntary conduct, make themselves morally liable to being killed or injured. In such cases, the relevant moral rights are forfeited and not infringed.

Alternatively, killing or injuring some may be the lesser evil, necessary to prevent greater harm to others. More precisely, killing or injuring some is morally permissible as a necessary means of preventing far greater harm to others or as a necessary side-effect of preventing substantially greater harm to others. This qualification reflects the view that, other things equal, intentionally doing harm is morally much worse (or harder to justify) than unintentionally doing harm which is, in turn, morally substantially worse (or harder to justify) than allowing harm. In such cases, the relevant moral rights are overridden and therefore justifiably infringed.

I propose that it is *jus in bello* proportionate to unintentionally kill or harm civilians as an unintended side-effect of achieving a military advantage that will prevent (or enable one to prevent) substantially greater harm to one's own forces or civilians⁹ over the remainder of the conflict.¹⁰ In these cases, the moral rights of the civilians harmed are permissibly overridden to protect the moral rights of others. On this view, the moral value of a military advantage lies in the harm to soldiers and civilians that it prevents. This account applies symmetrically to all sides of a conflict, independently of their war aims, yet identifies a morally compelling explanation for when military advantage justifies civilian losses.

Let me illustrate the account by applying it to a series of skeletal cases. To keep things simple, these cases will assume an international armed conflict between state armed forces. At the end of the chapter I will apply the underlying framework to targeted killing operations within a non-international armed conflict between state armed forces and non-state armed groups.

The first case is the simplest:

⁹ For the purposes of this chapter, "one's own civilians" includes all civilians that an armed force seeks to protect through the use of military force, whether fellow citizens, allies, occupied populations, or intended beneficiaries of humanitarian intervention. As a result, the civilians an attack may harm and the civilians an attack may protect are sometimes members of the same political community.

¹⁰ The last qualification is important because, while attacks on opposing forces sometimes prevent imminent harm to attacking forces or to civilians, most attacks on military equipment or strategic locations prevent or avoid such harm only over the remainder of the conflict.

Simple Prevention: State A is at war with State B. If State A destroys State B's long-range missiles, some of State B's civilians will be killed as a side-effect. However, if State A does not destroy State B's long-range missiles, many more of State A's soldiers and civilians will be killed by those missiles over the remainder of the conflict.

A strike on the long-range missiles directly prevents State B from harming State A's soldiers and civilians (hence *Simple Prevention*). On the view under consideration, this operation seems proportionate, since the immediate harm it inflicts is substantially less than the future harm it prevents.

It might seem that the preceding case calls for a simple comparison between the consequences of action and the consequences of inaction, but that is not quite correct. The case assumes that if the operation is not launched then State A will not simply surrender but will continue to resist State B by other means. However, State A's alternative means of resistance will not prevent State B from launching its long-range missiles. In this sense, the view under consideration takes it for granted that parties will pursue their war aims and compares only the harms inflicted and prevented by particular military operations.

Now let us add some facts to the first case:

Triple Prevention: If State A destroys State B's anti-aircraft missiles, some of State B's civilians will be killed as a side-effect. However, if State A does not first destroy State B's anti-aircraft missiles then State A's air forces will be shot down before they can destroy State B's long-range missiles.

A strike on the anti-aircraft missiles will not directly prevent State B from harming State A's soldiers and civilians. Instead, the strike will prevent State B from preventing State A from preventing State B from harming State A's soldiers and civilians (hence *Triple Prevention*).

On my view, *jus in bello* proportionality is satisfied only if the losses inflicted on State B's civilians by *both* the strike on the anti-aircraft missiles *and* the subsequent strike on the long-range missiles are substantially less than the losses the long-range missiles would inflict on State A's soldiers and civilians. Cases like *Triple Prevention* illustrate the general truth that the proportionality of an attack depends not only on the harm it prevents but also on the harm it enables additional attacks to prevent. For example, "[i]f, in order to prevent the enemy's army from advancing, planners decide to destroy all the bridges that span a river, . . . each driver or pilot may judge that his own action is disproportionate, [but] the operation as a whole may meet the proportionality requirement."¹¹

¹¹ Jean-François Quéguiner, *Precautions under the law governing the conduct of hostilities*, 88 INT'L REV. RED CROSS 804 (2006).

Now let us change the facts of the second case slightly:

Timely Prevention: If State A destroys State B's anti-aircraft missiles, some of State B's civilians will be killed as a side-effect. However, if State A does not first destroy State B's anti-aircraft missiles then it will take longer for State A's air forces to destroy State B's long-range missiles; these long-range missiles will kill many of State A's civilians before they are destroyed.

In this case, State A can prevent State B from using its long-range missiles to kill State A's soldiers and civilians without first destroying State B's anti-aircraft missiles. However, fewer of State A's civilians will be killed if State A first destroys State B's anti-aircraft missiles. Evidently, *Timely Prevention* is just a special case of *Triple Prevention*, in which we compare the losses inflicted by the attack on the anti-aircraft missiles with the losses prevented by destroying the long-range missiles sooner rather than later.

Now consider the following variation on the previous cases:

Costly Triple Prevention: If State A destroys State B's anti-aircraft missiles, some of State B's civilians will be killed as a side-effect. However, if State A does not first destroy State B's anti-aircraft missiles then State A's air forces will suffer substantial losses before destroying State B's long-range missiles.

In this case, State A can prevent State B from using its long-range missiles to kill State A's soldiers and civilians without first destroying State B's anti-aircraft missiles. However, fewer of State A's soldiers will be killed trying to prevent State B from using its long-range missiles if State A first destroys State B's anti-aircraft missiles. It follows that an attack on the anti-aircraft missiles is justified only if the losses *those* missiles would inflict on attacking forces are substantially greater than the civilian losses the attack itself would inflict.

Arguably, this case is not governed by *jus in bello* proportionality, but by *jus in bello* necessity as understood in Chapter 7. On this interpretation, the true military advantage in this case is the destruction of the long-range missiles, and attacking forces must choose between one means of achieving that advantage that involves greater harm to civilians but less harm to attacking forces and another means of achieving the same advantage that involves greater harm to attacking forces but less harm to civilians.

Alternatively, we could say that there are two distinct military advantages in this case. The value of destroying the long-range missiles lies in the losses this will prevent, assuming that the attacking force will not surrender and would otherwise pursue victory by other means. The value of destroying the anti-aircraft missiles lies in the losses that this will prevent, assuming that the attacking force will, and could proportionately, attack the long-range missiles. On this view, *Costly Triple Prevention* simply combines

two cases of *Simple Prevention*, one nested within the other. As described in Chapter 7, the proper classification of such cases is conceptually interesting but makes no practical difference. Both necessity and proportionality are best understood in terms of the moral asymmetry between doing harm and allowing harm, and will permit or prohibit the same actions.

As a general matter, attacks on opposing combatants, military equipment, and military facilities will share the same causal structure as the four preceding cases. Such attacks either prevent the object of attack from inflicting future harm, prevent the object of attack from preventing the attacking force from preventing some further object of attack from inflicting future harm, prevent the object of attack from delaying the attacking force from preventing some further object of attack from inflicting future harm, or prevent the object of attack from inflicting harm on the attacking force or its civilian population as it seeks to prevent some further object of attack from inflicting future harm. Such attacks are proportionate if they prevent (or enable the prevention of) substantially greater harm than they (or the sequence of attacks of which they are a necessary part) inflict.

Attacks that aim to secure strategic locations introduce an additional causal step between the initial attack and the prevention of harm. Such attacks either indirectly prevent harm by denying such locations to the opposing force or enable the attacking force to prevent harm in subsequent operations. Though the causal structure of such cases is more complex, their moral structure is the same: such attacks are proportionate only if they indirectly prevent (or indirectly enable the prevention of) substantially greater harm than they (or the sequence of attacks of which they are a necessary part) inflict.

Strategies and Tactics

We can now see where, on my account, the value of a military advantage lies. Let X represent the total losses that one's own forces and civilians will suffer in the remainder of the conflict if a military advantage (say, destroying a munitions factory or killing an insurgent leader) is not achieved. Let Y represent the total losses that one's own forces and civilians will suffer in the remainder of the conflict if that military advantage is achieved. The value of the military advantage is the difference between X and Y , that is, the total overall losses prevented by achieving the advantage. Let Z be the losses that one would unintentionally inflict on opposing civilians in pursuit of the advantage. On my view, an attack is proportionate just in case the difference between X and Y is substantially greater than Z .

I have described the value of a military advantage in terms of preventing losses to attacking forces and to civilians. On one hand, some may argue that only avoided losses to attacking forces should count toward the value of a *military* advantage. This view seems clearly false. In wars of

national self-defense, humanitarian intervention, military occupation, and population-centric counter-insurgency the protection of civilian populations is (or morally should be) a primary war aim. It must count in favor of the proportionality of an attack that it will prevent the opposing force from inflicting substantially greater harm on civilians.

On the other hand, some may observe that it is hardly ever the primary aim of an attacking force to reduce or avoid losses to its members. The primary aim is almost always something else, such as the protection of a civilian population or of the territorial integrity or political independence of a state. Avoiding or reducing losses to one's own forces is at most a moral or pragmatic side-constraint on the pursuit of some broader aim. As we have seen, Hurka and McMahan argue that we should derive the value of a military advantage from the degree of its contribution to the achievement of that broader war aim.

Here we must distinguish between the *concept* of a military advantage and the *value* of a military advantage. The *concept* of a military advantage may refer to the contribution of an attack to a party's war aims. But it does not follow that the *value* of a military advantage must be understood in terms of its contribution a war aim. For one thing, how could we divide up the value of achieving a war aim among contributory military advantages? By definition, a *contribution* to a war aim is either a necessary contribution to its achievement or a necessary member of a sufficient set of contributions to its achievement.¹² There is no such thing as a superfluous contribution. So if the value of a military advantage is determined by its contribution to a war aim then it follows that the value of every military advantage is equivalent to the value of achieving the broader war aim. Obviously, this is not an attractive position. Among other things, this view eviscerates the notion that *jus in bello* proportionality compares the harms and benefits of particular attacks rather than those of the war as a whole.

Evidently, necessary members of a sufficient set of contributions do not inherit the total value of the outcome. Instead, the value of a necessary member of a sufficient set of contributions can only be determined by comparing the value of the set with the value of alternative sufficient sets of

¹² Of course, an advantage might contribute to achieving a war aim sooner rather than later. However, since the main value of shortening a war lies in reducing losses, this nuance does not change the analysis.

Jeff McMahan pointed out to me that a series of attacks could over-determine the achievement of a war aim. In such cases, each attack causally contributes to success even though success is not counter-factually dependent on any individual attack. In response, I would claim that an act that infringes rights can only be justified by its results if those results are counter-factually dependent on the act. If I am right, then attacks which over-determine the achievement of a war aim may each be disproportionate in the fact-relative sense. Of course, since the results of an attack are always difficult to predict, it may be permissible in the evidence-relative sense to launch a series of attacks some of which will prove impermissible in the fact-relative sense. I suspect that many actual cases of over-determination in war are best understood along these lines.

contributions. Suppose that there are N sufficient sets of contributions to producing G . The value of any particular contribution C to set S_1 should be determined by comparing the value of S_1 with the value of sets S_2 through S_N . The value of C is not G (the good to which C contributes) but rather the difference in value between S_1 and S_2 through S_N . On this view, the (gross) costs of C are proportionate if they are outweighed by the (net) benefit of choosing S_1 over S_2 through S_N .

In the context of armed conflict, the value of a military advantage is the difference in value between a broader strategy that requires achieving that advantage and alternative strategies that do not. Harm to civilians inflicted in pursuit of a military advantage is proportionate only if any alternative strategy would allow substantially greater harm to civilians and to the attacking force. More precisely, if an attack that is part of strategy S_1 inflicts degree of harm H on opposing civilians, then the attack is proportionate only if S_2 through S_N would occasion more harm to one's own forces and civilians than S_1 , and if the net difference in overall harm is substantially greater than H .¹³

Why should the proportionality of a military operation depend on a comparison between the strategy of which it is a part and alternative strategies in which it is not a part? Such a comparison is unavoidable because the harms an attack will prevent are just the harms that will occur if the attack is not carried out, which in turn depend on what the attacking force would or could do instead. Again, *jus in bello* proportionality assumes that parties will pursue their aims by one strategy or another. The harm an attack prevents is just the harm that the attacking force would suffer if it pursued an alternative strategy.

Now, it is sometimes (through rarely) the case that a military advantage makes a necessary contribution to a war aim, such that there is no sufficient set of contributions that does not include it as a necessary member. If foregoing a military advantage would make victory impossible, then an attack that achieves that advantage is proportionate if the moral value of achieving the war aim outweighs the expected harm to civilians. Such cases could occur in broader conflicts or, in principle, in the context of single defensive strikes against discrete threats. In such cases, and only in such cases, *jus in bello* proportionality collapses into *jus ad bellum* proportionality. In such cases, and only in such cases, only attacking forces who fight for a just and lawful cause can fully satisfy *jus in bello* proportionality.

Let me take a moment to distinguish my view from one that McMahan effectively criticizes, namely the view that *jus in bello* proportionality compares harms inflicted on civilians with harms avoided to combatants *in the very same engagement*. On this view, just combatants may inflict harm on civilians if such harm is a necessary and proportionate side-effect of using

¹³ Of course, S_1 may itself violate *jus ad bellum* proportionality, if the moral value of its aim is outweighed by the harm to civilians expected in pursuit of that aim.

defensive force against unjust combatants. This view could at most explain the permissibility of *unit self-defense*, that is, of force used to repel a direct attack on particular combatants. However, as McMahan observes, this view would preclude all offensive operations, since if a party refrains from offensive operations then the necessity to use defensive force on behalf of its members often will not arise. In particular, McMahan argues that this view would preclude humanitarian military interventions, since in these interventions attacking forces will inflict losses on civilians that are not strictly necessary to protect their members since the attacking forces could simply not intervene. In addition, it is not clear how this view would apply to targeted killing operations in which the attacking force is in no immediate danger.

My view is not subject to these objections. On my view, we hold constant the war aim of the attacking force (national self-defense, civilian protection, and such like) and ask whether an attack is a necessary part of a broader strategy for achieving that aim. Offensive operations are proportionate, including as part of a humanitarian intervention, only if the immediate losses they inflict are substantially less than the future losses they (or the sequence of attacks of which they are a necessary part) will prevent, assuming always that the attacking force will pursue its war aims.

Moral Standards and Decision Procedures

Hopefully, my view is appealing on its own terms. Interestingly, my view derives further support from an appealing recent proposal, combined with some plausible assumptions. In unpublished work, both Avishai Margalit and Jeremy Waldron have (it seems independently) arrived at the following proposal:

M/W: Attacking forces may inflict losses on civilians in pursuit of a military advantage only if those losses do not exceed the losses they would accept to their own members in pursuit of that military advantage.

For example, *M/W* would direct a commander not to order a missile strike against a militia leader that would unintentionally kill three civilians unless the commander would be willing to lose three soldiers to achieve the same goal. The commander might consider whether he or she would order a ground assault that would claim the lives of three soldiers, or whether he or she would order the strike even if the strike would itself kill three soldiers in so-called “friendly fire”. Of course, the actual missile strike presumably will not put any actual soldiers in harm’s way, so it is important to make clear that *M/W* is a hypothetical test like the Golden Rule. The relevant question is what losses the attacking force would be willing to incur to achieve the advantage, not what losses they will in fact incur in its pursuit.

The attraction of *M/W* is that it demands that attacking forces take equally seriously losses to their own members and losses to civilians. Since attacking forces generally take the former far more seriously than the latter, and since the latter is at least as morally important as the former, *M/W* serves a valuable corrective function in moral deliberation. Moreover, since military commanders are presumably competent to determine how many of their own soldiers they are willing to lose to achieve a given military advantage, *M/W* provides military commanders with an administrable decision procedure.

Importantly, *M/W* is *only* a decision procedure and not a moral standard. For one thing, much like the Golden Rule, *M/W* depends for its application on the subjective values of the applicant. The less value a commander places on the lives of his or her own forces, the less value *M/W* requires him or her to place on the lives of civilians. *M/W* tells us only to place equal value on the lives of our own forces and the lives of civilians. But *M/W* does not tell us how much value to place on the lives of anyone. To convert *M/W* into a moral standard we need some way to objectively determine how much harm a commander should allow to befall his or her own forces in pursuit of a given military advantage.

Now, a *rational* commander exclusively concerned with the lives and welfare of his or her own forces will accept the following claim:

Own: Attacking forces may accept losses to their members in pursuit of a military advantage only if achieving that military advantage would prevent greater losses to their members over the remainder of the conflict.

In other words, attacking forces may pursue a military advantage only if the immediate losses to their own members in pursuit of that military advantage will be less than the additional losses to their own members if the military advantage is not achieved. Since there is no obvious moral asymmetry between occasioning immediate losses to one's own forces and occasioning future losses to one's own forces, only the numbers count.

M/W tells *moral* commanders to treat civilians as they treat their own soldiers. *Own* tells *rational* commanders how to treat their own soldiers. Combining *M/W* with *Own*, we begin to see how a rational *and* moral commander should treat civilians:

A: Attacking forces may inflict losses on civilians in pursuit of a military advantage only if achieving that military advantage would prevent greater losses to their members over the remainder of the conflict.

Obviously, the problem with the preceding claim is that it ignores the fact that sometimes achieving a military advantage will reduce overall losses to civilians as well as to one's own forces. Since reducing losses to civilians is

clearly morally important, a rational and moral commander will accept the following refinement:

B: Attacking forces may inflict losses on civilians in pursuit of a military advantage only if achieving that military advantage would prevent greater losses to their members *and to civilians* over the remainder of the conflict.

The problem with this claim is that it ignores the moral asymmetry between doing harm and allowing harm. In general, it is permissible to unintentionally do harm only as a side-effect of preventing substantially greater harm. Of course, once we incorporate the relevant moral asymmetry, we arrive at my position:

C: Attacking forces may inflict losses on civilians in pursuit of a military advantage only if achieving that military advantage would prevent *substantially* greater losses to their members and to civilians over the remainder of the conflict.

Thus, my view may stand on its own or on the strength of *M/W* and plausible assumptions regarding rationality and morality.

Fellow Citizens and Foreign Civilians

I have argued that an attack is proportionate if the harm it prevents is substantially greater than the harm it inflicts. In internal armed conflicts and humanitarian interventions, attacks often inflict and prevent harm to members of the same civilian population. However, in international armed conflicts attacks typically inflict harm on foreign civilians while preventing harm to members of the attacking force or its civilian population. Some scholars argue that, in such cases, combatants have *special duties* to protect their comrades and fellow citizens. These special duties may be associative obligations based on shared nationality or role-responsibilities voluntarily undertaken. Importantly, these scholars argue that these special duties make it morally permissible for combatants to inflict harm on foreign civilians to prevent equal or lesser harm from befalling the attacking force or its civilian population.¹⁴ On these views, the special duties of combatants offset or override the general moral asymmetry between doing and allowing harm.

There are many difficulties with such views.¹⁵ Certainly, combatants generally bear special duties to protect their comrades and civilians from harm. These duties may be very *stringent*, in the sense that combatants may be morally required to assume great risks and bear great burdens to

¹⁴ See, e.g., Hurka, at 63-64.

¹⁵ See, e.g., Adil Ahmad Haque, *Killing in the Fog of War*, 86 SOUTHERN CALIFORNIA LAW REVIEW 63, 86-88 (2012).

discharge these duties. However, it does not follow that combatants are morally permitted to impose great risks or burdens on foreign civilians to discharge their own duties.

Consider a case of promissory obligation. If I promise a friend to meet at a particular place and time then I assume a special duty to do so. This special duty is only mildly stringent: I may be morally required to pay \$20 for a taxi if necessary to get there on time, but perhaps not to pay \$200. Yet I am not morally permitted to steal \$20 from a bystander, or to cause a bystander to lose \$20, as a means or side-effect of getting there on time. It is, after all, *my* duty to get there on time, and the burden of discharging it should fall on me alone. I cannot force bystanders to bear those burdens any more than I can transfer my duty to them. Of course, if failing to meet on time would somehow result in harm to my friend then I may be morally permitted to inflict a lesser harm on a bystander as a side-effect of preventing a substantially greater harm to my friend. But at this stage my special duty drops out of the picture and the general morality of doing and allowing harm takes over.

Similarly, combatants may be morally required to lay down their own lives to discharge their special duties to protect their comrades and fellow citizens. Yet combatants are not morally permitted to lay down the lives of foreign civilians simply to discharge their own duties. Combatants are permitted to kill foreign civilians only according to the general morality of doing and allowing harm. To allow their special duties to affect proportionality would be to force foreign civilians to bear the burdens of discharging someone else's duties.

Here is another way of seeing that special duties cannot affect the moral permissibility of killing. Suppose an innocent stranger is about to be killed and asks me to defend her. She retains her right to kill in self-defense and may transfer that right to me. However, my general duty to defend her is not very stringent and I may decline to defend her if doing so would place me at great risk. Now suppose that an innocent civilian is about to be killed and asks a soldier in her state's armed forces to defend her. The soldier's special duty to defend her is very stringent, as we have seen, such that the soldier must defend her even at great personal risk. But notice that *what* the soldier is required to do (the *content* of the duty) is to exercise the civilian's right of self-defense on her behalf. It follows that if the civilian may not inflict a greater harm on a foreign civilian as a side-effect of defending herself from a lesser harm then the soldier may not do so on her behalf. After all, no one can transfer a right that she does not possess. It follows that the special duties of soldiers to their civilian population do not confer on soldiers any special rights to inflict greater harm on foreign civilians as a side-effect of preventing lesser harm to their comrades or fellow citizens.

Since special duties do not affect the morality of doing and allowing harm, we can propose the following rule of thumb modeled on *M/W*:

Ours=Theirs: Attacking forces may inflict losses on civilians in pursuit of a military advantage only if those losses do not exceed the losses they would accept to their own civilians in pursuit of that military advantage.

Ours=Theirs instructs attacking forces not to launch an attack that inflicts losses on foreign civilians unless they would inflict comparable losses on their own civilians in pursuit of the same military advantage. Like *M/W*, *Ours=Theirs* is only a decision procedure, not a moral standard. However, if an attacking force may inflict losses on its own civilians only to prevent substantially greater losses to its own civilians, and if special duties do not affect the morality of doing and allowing harm, then *Ours=Theirs* entails the moral standard that I have proposed.

Symmetry and Independence

One virtue of my approach is that it allows both participants and observers to evaluate the proportionality of particular military operations without reference to the justice or legality of the broader conflict. On my view, *jus in bello* proportionality assumes that both sides will pursue their war aims and asks whether a particular attack prevents substantially greater harm to the soldiers and civilians of one side than it inflicts on the civilian population of the other. The broader permissibility of pursuing particular war aims through military force, however restrained, is a matter of *jus ad bellum* and not *jus in bello*.

Certainly, the law requires both the independence of *jus in bello* proportionality from *jus ad bellum* considerations and the symmetrical application of *jus in bello* proportionality to all sides of a conflict. A legal rule that flatly prohibits harming civilians in pursuit of an unjust cause will be ignored both by combatants who believe their cause is just and by combatants who feel compelled to fight for a cause they believe is unjust. The law assumes that warring parties will seek to pursue their war aims at the least possible cost to their own soldiers and civilians. *Jus ad bellum* proportionality constrains a party's pursuit of its war aims by comparing the value of their achievement with the cost of their pursuit to opposing civilians. In contrast, *jus in bello* proportionality constrains a party's attempts to reduce its own losses in pursuit of its aims by comparing the future harm an operation prevents with the immediate harm it inflicts.

Not everyone will see the symmetrical application of my proposal as a virtue. Recall that Hurka and McMahan argue that a military advantage that does not contribute to the achievement of a just cause has no moral value that could justify the moral disvalue of harm to civilians. They conclude that any harm inflicted on civilians by unjust combatants in pursuit of an unjust cause is disproportionate from a moral point of view.

Hurka and McMahan are clearly right that attacks by unjust combatants are hardly ever *fully* morally justified, since they generally kill and injure just combatants, make the achievement of unjust causes more likely, and are in one sense unnecessary if the opposing party would agree to a ceasefire and ultimately to a just peace.¹⁶ But attacks by unjust combatants can be *partially* morally justified to the extent that they prevent harm to others. Since almost all civilians retain their moral right not to be killed, there is almost always strong moral reason to prevent civilians from being killed. Similarly, since even unjust combatants hardly ever deserve to die, there is almost always some moral reason to prevent them from being killed. An attack is morally better to the extent that it is supported by such reasons and morally worse to the extent that it is not. Since proportionality is in part a function of such reasons, it is almost always *morally worse* for unjust combatants to launch disproportionate attacks than for them to launch proportionate attacks. It follows that even combatants who knowingly fight without a just cause (perhaps under duress) have significant moral reasons to comply with *jus in bello* proportionality. In doing so they will seldom act permissibly, but they will almost always act less wrongfully than they would otherwise.

As we saw in Chapter 5, the moral function of the law of armed conflict is to help all combatants, just and unjust alike, conform to the moral reasons that apply to them. For unjust combatants, full conformity is generally impossible and partial conformity is the most they can achieve so long as they continue to fight at all. But given the moral stakes of armed conflict, even partial conformity is most welcome. If unjust combatants stop killing civilians except as a side-effect of preventing substantially more civilians or combatants from being killed then armed conflict will be much less unjust than if they are not so constrained.

Collective Goods

One virtue of my approach is that it only requires comparisons between commensurable losses, namely death and injury.¹⁷ What, then, of collective goods such as national self-determination? Since most scholars assume that such collective goods perform a crucial role in *jus ad bellum*, must they not perform a crucial role in *jus in bello*?

¹⁶ Of course, generally only the most senior political and military leaders can negotiate a ceasefire. Individual unjust combatants often must choose between harming just combatants and allowing harm to their own civilians. In some such cases, individual unjust combatants may be fully morally justified in defending their own civilians.

¹⁷ Attacks that damage civilian property are proportionate only if they prevent greater damage to military or civilian property, or death or injury to soldiers or civilians. Conversely, death and injury to civilians can never be justified solely to prevent damage to property.

As we saw in Chapter 3, collective goods play a different role in justifying the conduct of hostilities than is generally supposed. National self-determination, for example, is not a goal that permissibly may be pursued by means of killing combatants. Rather, the value of self-determination explains why soldiers have no duty to retreat or surrender when their efforts to preserve their self-determination meet with armed resistance.

Similarly, self-determination is not a good that must be weighed against the evil of inflicting losses on civilians. The primary justification for inflicting losses on civilians is that one would otherwise allow substantially greater losses to one's own forces or to other civilians. Self-determination explains why one need not retreat or surrender in order to avoid the need to choose between inflicting and allowing such losses.

Let me put the point a different way. One could fairly ask: how can an attack that is proportionate in the sense I have defended be permissible, given that one could avoid losses to anyone by simply retreating or surrendering? In some cases, involving genocidal aggression for example, losses cannot be avoided because the opposing force seeks to impose such losses as an end in itself. In these cases, the threatened party is put to a forced choice in a *descriptive* sense: there is no course of action that does not involve inflicting or allowing serious losses.

However, in many cases, involving so-called "conditional" or "lesser" aggression, the opposing force will only inflict losses if it is denied its political objective. In these cases, the threatened party is not morally required to retreat or surrender in the face of conditional aggression and cede control over their territory and population. They may resist, and if their resistance is or will be met with force they may defend themselves with force. If their use of defensive force would inflict losses on civilians, then the losses they inflict must be substantially less than the losses they prevent or avoid. But they need not avoid inflicting otherwise proportionate losses by surrendering their self-determination. In these cases the threatened party is put to a forced choice in a *moralized* sense: there is a course of action that does not involve inflicting or allowing serious losses, but it is not a course of action that morality requires. Of course, a party may choose to sacrifice its self-determination rather than inflict and occasion losses to civilians, but such a sacrifice is supererogatory and not morally required.

Excessiveness and Proportionality

The law of armed conflict prohibits attacks that are expected to inflict harm on civilians that is 'excessive' in relation to the anticipated military objective. Read narrowly, the term "excessive" simply means "too great" and leaves entirely open the standard by which to determine when harm to civilians is too great in relation to anticipated military advantage. However, it is not facially implausible to think that the term "excessive" indicates that the

harms of the attack must *exceed* the benefits of the attack in order to render the attack unlawful. Combined with my account of the value of a military advantage, this understanding of “excessive” would entail that an attack is lawful even if it prevents no more harm than it inflicts.

This interpretation of the law is linguistically plausible but morally defective. Since killing is substantially worse than letting die, it is proportionate to inflict unintentional harm on civilians only if necessary to avoid substantially greater harm to other civilians or attacking forces. For example, if an attack on an opposing military unit, either on one’s own territory or during a humanitarian intervention, would prevent that unit from killing 100 civilians but would also itself kill 100 civilians (of the same political community, let us suppose) then the attack is morally impermissible.

The morally sound view, and an equally plausible interpretation of the legal rule, is that losses to civilians are “too great” if their moral disvalue exceeds or outweighs its moral value. On this view, moral value includes both consequential and deontic value. Since it is morally worse to inflict losses than to allow losses, losses inflicted on civilians are excessive unless they occur as a side-effect of avoiding substantially greater losses to attacking forces and civilians.

As we saw at the beginning of this chapter, proportionality judgments are somewhat imprecise. For example, it is seldom the case that an attack that is proportionate if it harms N civilians becomes disproportionate if it harms $N+1$ civilians. Most scholars suppose that proportionality judgments are imprecise because they compare distinct and incommensurable values, namely harm to civilians and military advantage. Of course, I have argued that the value of a military advantage lies in the harm prevented by its achievement. Thus, on my account, proportionality judgments compare fully commensurable values, namely harms to some and harms to others. On my view, proportionality judgments are imprecise because the moral asymmetry between doing and allowing harm is imprecise. When we say that killing is substantially worse than letting die we do not mean that killing is *precisely* X times worse than letting die, where X is an exact number like 3.14. Rather, we mean that killing is *at least* X times worse than letting die, where X is some substantial figure. Moral philosophers have not devoted much effort to assigning even an imprecise value to X , so it would be presumptuous to assume that a precise value could be determined.

Concrete and Direct Advantages

Importantly, the law of armed conflict weighs harms to civilians only against military advantages that are ‘concrete and direct’. Since my account of the value of a military advantage assigns critical importance to the impact of an attack on the total losses suffered over the remainder of a conflict, it is important to show that my account is consistent with existing law.

It should first be admitted that the meaning of the terms “concrete and direct” are not well-settled in international law. For example, the ICRC comments that ‘the advantage concerned should be substantial and relatively close, and . . . advantages which are hardly perceptible and those which would only appear in the long term should be disregarded’.¹⁸ In contrast, the ICC Elements of Crimes states that, under the parallel provision of the Rome Statute, ‘[t]he expression “concrete and direct overall military advantage” refers to a military advantage that is foreseeable by the perpetrator at the relevant time. Such advantage may or may not be temporally or geographically related to the object of the attack.’¹⁹ Of course, advantages that are temporally or geographically remote are often difficult to reliably foresee, in part because the necessary causal chain may be broken at more points. But proximity is at best a proxy, not a substitute, for foreseeability.

The moral relevance of foreseeability is obvious (and well-explained in Chapter 6): if an outcome is unlikely to occur then its pursuit cannot justify inflicting harm on civilians. The law will serve soldiers well by excluding low-probability advantages from their proportionality calculations. The moral relevance of temporal and geographical proximity is harder to understand. For example, the military advantage of destroying a munitions factory generally will not be realized until existing munitions are depleted and new munitions needed. Yet the military advantage of destroying a munitions factory is paradigmatically legitimate and (in principle) capable of justifying the unintentional infliction of civilian losses.

More broadly, we now see that any morally defensible legal norm cannot compare the immediate harm inflicted by an attack only with the harm directly prevented by that attack (as in *Simple Prevention*). The proportionality of an attack must also consider future harms prevented by later attacks that the first attack enables (as in *Triple Prevention* and *Indirect Simple Prevention*) or that are enabled by attacks that the first attack enables (as in *Indirect Triple Prevention*). The drafters of the Rome Statute were therefore wise to compare civilian losses to the ‘concrete and direct *overall*’ military advantage produced by an attack. The proportionality of an attack often depends on the proportionality of the sequence of attacks of which it is a necessary part.²⁰

¹⁸ Protocol I Commentary, para. 2209. However, as Rogers points out, the Commentary explanation of the words ‘concrete and direct’ does not accord with ordinary English usage. The meaning of the word ‘concrete’ is not ‘predicted’ but ‘specific’ or ‘perceptible’ (rather than general) and the meaning of the word ‘direct’ is not ‘relatively close’ but ‘without intervening condition of agency’. ROGERS, LAW ON THE BATTLEFIELD 59-60.

¹⁹ U.N. Preparatory Comm'n for the Int'l Crim. Court, Finalized Draft Text of the Elements of Crimes, fn. 36, U.N. Doc. PCNICC/2000/1/ADD.2 (November 2, 2000).

²⁰ Several leading states (the United Kingdom, Belgium, Canada, the Federal Republic of Germany, the Netherlands, Italy, Spain, and the United States) took the position that the reference to ‘direct military advantage anticipated’ from an attack was intended to refer to the advantages anticipated from the attack considered as a whole, and not just from

There are also categorical restrictions on the proportionality calculus. For example, the ICRC comments that ‘there can be no question of creating conditions conducive to surrender by means of attacks which incidentally harm the civilian population. A military advantage can only consist in ground gained and in annihilating or weakening the enemy armed forces.’²¹ The proposed restriction seems too strong. It is true that the political impact of unintentionally killing civilians cannot contribute to the proportionality of an attack. As we saw in Chapter 2, the causal consequences of killing civilians cannot even partially justify their deaths. However, there is no obvious reason why the political impact of killing combatants or destroying military objectives cannot contribute to the proportionality of an attack that also harms civilians. For example, if destroying legitimate military targets will both degrade the military capacity of the opposing party and deter the opposing party from launching future attacks then it seems that both consequences can help justify any unintended harm to civilians. Put another way, deterrence can contribute to proportionality so long as it results from destroying military targets rather than from harming civilians

To sum up, the best view seems to be that military advantages are ‘concrete and direct’ only if (a) their value is sufficiently great to justify harming civilians; (b) they are reasonably foreseeable at the time of attack; (c) they result from the attack itself or from the sequence of attacks of which the attack is a necessary part; and (d) they result from damage to military targets rather than from harm to civilians.

Uncertainty

Importantly, the law does not call for a comparison between the actual civilian losses inflicted by an attack and the actual military advantage the attack achieves. Instead, the law calls for a comparison between the civilian losses *expected* and the military advantage *anticipated*. Since the law aims to guide conduct, it must provide soldiers with a decision procedure that can be followed under conditions of epistemic uncertainty. Soldiers can seldom know the exact consequences of an attack with certainty. Instead, an attack might have various outcomes, some more likely than others. If the law is to perform its guidance function it must tell soldiers how to compare the various possible consequences of their actions.

One might think that an attack is permissible if the *expected* harm to civilians (that is, the average of the possible harms to civilians discounted by their respective probabilities) is substantially less than the *expected* military advantage (that is, the value of the military advantage discounted by the

isolated or particular parts of the attacks during the drafting of Protocol I, and there were no recorded objections. See H.S. LEVIE, PROTECTION OF WAR VICTIMS: PROTOCOL I TO THE 1949 GENEVA CONVENTION 165, 169, 171, 331-132, 334 & 336-137 (1980).

²¹ Protocol I Commentary, para. 2218.

probability of its achievement). In Chapter 6, we described this concept as *e-proportionality*. When an attack will probably not harm civilians, then generally it is permissible if it is e-proportionate. In general, it is morally permissible to take a small risk of immediately harming civilians for the sake of a substantially greater chance of preventing comparable harm to others or a comparable chance of preventing substantially greater harm to others.²² However, as we saw in Chapter 6, an attack that will probably harm civilians may not be permissible even if it is e-proportionate because the *expectably-best* action is not always permissible in the evidence-relative sense.

Suppose that a targeted killing operation directed at a particular low-level insurgent will probably (more likely than not) kill several civilians and prevent little or no harm to soldiers. However, there is a small chance that this low-level insurgent will one day develop into a senior leader who will make a necessary contribution to the killing of many soldiers. Suppose that the moral gravity of this remote possibility is sufficient to make the expected military advantage of killing the insurgent substantially greater (though not far greater) than the expected harm to civilians. On these facts, the strike is still morally impermissible in the evidence-relative sense because you cannot reasonably believe that the attack will prevent substantially more harm than it will inflict.

In general, an action is permissible in the evidence-relative sense only if you have undefeated reason to believe (a) that the action will infringe no rights, (b) that a justifying circumstance exists, or (c) that a justifying result will occur. Conversely, an action is impermissible in the evidence-relative sense if you have decisive reason to believe (a) that the action will infringe rights, (b) that no justifying circumstance exists, and (c) that no justifying result will occur. It follows that you may not carry out an attack that you reasonably believe will harm civilians (an infringement of unforfeited rights) unless you reasonably believe that the attack will prevent substantially greater harm to friendly forces or friendly civilians (a justifying result). Put another way, attacking forces must ask two questions:

- (i) How many civilians will the attack probably kill?
- (ii) How many soldiers and civilians will the attack probably save?

An attack is permissible in the evidence-relative sense only if the answer to (ii) is substantially greater than the answer to (i).

Often, there will be no exact number of civilians that an attack will probably kill or of soldiers and civilians that an attack will probably save. To keep things simple, suppose there is a 40% chance that an attack will kill fewer than 10, a 40% chance that it will kill between ten and fifteen, and a 20% chance that it will kill between 15 and 20. Similarly, suppose there is a

²² See HUGO GROTIUS, THE LAW OF WAR AND PEACE 601 (1625, 1925): '[B]eware of what happens, and what we foresee may happen, beyond our purpose, unless the good which our action has in view is much greater than the evil which is feared, or, unless the good and the evil balance, the hope of good is much better than the fear of the evil'.

40% chance that an attack will save fewer than 20, a 40% chance that it will save between 20 and 30, and a 20% chance that it will between 30 and 40.²³

	Kill	Save
40%	$K < 10$	$S < 20$
40%	$10 < K < 15$	$20 < S < 30$
20%	$15 < K < 20$	$30 < S < 40$
Max	$K < 15$	$S < 30$
Min	$10 < K$	$20 < S$

The attacking force cannot reasonably believe that it will kill fewer than 10 or save more than 30. It follows that if killing is more than three times worse than letting die then the contemplated attack is epistemically impermissible. More generally, we can propose

Min-Max: If an attack will probably kill *at least* N civilians, and if killing is X times worse than letting die, then the attack is epistemically permissible only if it will probably prevent *at most* $X(N)$ soldiers and civilians from being killed.

Min-Max compares the minimum number one can reasonably believe an attack will kill with the maximum number one can reasonably believe an attack will save. *Min-Max* is the correct minimum moral standard of epistemic permissibility, since it permits action based on reasonable beliefs that, if true, would make the action objectively permissible. However, there is good reason for the law to impose a more demanding standard.

To its credit, *Min-Max* prohibits pure wishful thinking on the part of attacking forces. However, given the potential for motivated reasoning in predictive judgments, attacking forces will often underestimate the harm they will inflict and overestimate the harm they will prevent. To offset this tendency, the law should demand that attacking forces instead compare the maximum number they can reasonably believe an attack will kill with the minimum number they can reasonably believe an attack will save. In the scenario above, the attacking force could (though need not) reasonably believe that the attack will kill 15 and save 20. It follows that if killing is even somewhat worse than letting die then the attack is epistemically impermissible. In general form, we can propose

Max-Min: If an attack will probably kill *at most* N civilians, and if killing is X times worse than letting die, then the attack is

²³ Importantly, the number killed and the number saved may vary independently of one another. I am using the same probabilities for illustrative purposes only.

epistemically permissible only if it will probably prevent *at least* $X(N)$ soldiers and civilians from being killed.

In principle, we could construct other permutations (call them *Max-Max* and *Min-Min*), but there is little to recommend them. There is no reason to believe that soldiers underestimate both the likelihood of killing and the likelihood of saving or overestimate both the likelihood of killing and the likelihood of saving. To the extent that soldiers engage in motivated reasoning they will underestimate the likelihood of killing and overestimate the likelihood of saving, the tendency that *Max-Min* is designed to control.

The law can fairly expect attacking forces to predict the harm an attack may inflict. At the very least, attacking forces generally know the blast radius of their own weapons and generally can determine the number and location of civilians in or near the target. In addition, attacking forces often can predict the effect of their weapons on nearby structures and identify at least some combustible or other hazardous materials in or near the target.

Nor is it unreasonable to expect attacking forces to predict the harm an attack may prevent. As we have seen, attacking forces must make such predictions simply in order to justify putting their own forces in immediate danger. Even a purely self-interested attacking force will only carry out attacks if the attack will probably reduce total harm to its members by preventing more harm over the remainder of the conflict than it occasions in the particular engagement. Such predictions are a necessary part of sound military decision-making.

In Chapter 5, I argued that reasonable belief that an individual is a combatant generally is not sufficient to epistemically permit lethal action given the moral risk involved. However, reasonable belief that an attack will prevent substantially more harm than it unintentionally inflicts generally is sufficient to epistemically permit the attack. In the target verification context, a higher level of certainty is often required to reflect the moral asymmetry between killing civilians and allowing soldiers to be killed. In the proportionality context, the moral asymmetry between killing and letting die is built into the substantive standard of proportionality itself. There is no need to build the same moral asymmetry into the relevant epistemic standard.²⁴

Finally, as we saw in Chapter 6, generally an attack that inflicts losses on civilians is not permissible if it will *probably* or *most likely* fail to achieve any military advantage, for example because a strike may not destroy its intended target. Such an attack could only be justified if the expected

²⁴ In addition, if an apparent combatant turns out to be a civilian then intentionally killing him or her is completely unjustified in the fact-relative sense. In contrast, if an attack on a verified military target prevents less, as much, or only slightly greater harm than it inflicts then generally it remains at least partially justified. It follows that the relative costs of error will seldom prove so uneven as to demand a very high standard of certainty.

military advantage is much greater (that is, between far greater and substantially greater) than the expected harm to civilians.

Superiors and Subordinates

Until now we have discussed the duties of attacking forces as collective agents without attending to the moral division of labor within such collectives. However, the moral and legal duties of commanders and subordinates differ in potentially significant ways.

Protocol I contains two distinct proportionality norms. The first norm requires ‘those who plan or decide upon an attack . . . [to] refrain from deciding to launch any attack’ that would inflict disproportionate harm on civilians.²⁵ This norm clearly applies only to those who design and order military operations. In contrast, the second norm states that ‘an attack shall be cancelled or suspended if it becomes apparent’ that the attack will inflict disproportionate harm on civilians.²⁶ The use of the passive voice suggests that this norm applies to all soldiers at all levels, including those who only carry out attacks. According to the ICRC, ‘[t]he rule set out here . . . applies not only to those planning or deciding upon attacks, but also and primarily, to those executing them’.²⁷

The introduction of the term “apparent” might suggest that in close cases it may be unlawful for a superior to order an attack but lawful for a subordinate to carry out an attack. On this interpretation, it is unlawful for a superior to order an attack if she has decisive reason to believe that it is disproportionate; but it is unlawful for a subordinate to carry out an attack only if he or she has strongly decisive reason, or perhaps conclusive reason, to believe that it is disproportionate. On reflection, however, this interpretation is difficult to accept. After all, the norm regarding cancellation and suspension of attack applies to superiors and subordinates alike. The proposed interpretation therefore entails that a superior may order an attack unless he or she has decisive reason to believe it is disproportionate but need only cancel or suspend the attack if he or she has strongly decisive or conclusive reason to believe it is disproportionate. It is hard to imagine what could justify imposing one epistemic standard on commanders before they order an attack and a different epistemic standard on the same commanders after the order is given but before the attack is carried out. Presumably, a quantum of evidence sufficient to rationally form a belief is also sufficient to rationally revise that belief.

A better view might be that new evidence regarding possible harm to civilians might give superiors decisive reason to believe that an attack is disproportionate but not give subordinates decisive reason to believe that an

²⁵ Protocol I, art. 57(2)(a)(iii).

²⁶ Protocol I, art. 57(2)(b).

²⁷ Protocol I Commentary, para. 2220. *See also* Quéguiner, at 803.

attack is disproportionate. After all, subordinates will not always know the value of the military advantage that an attack would yield; thus, they cannot always know when new information regarding possible harm to civilians will render the attack disproportionate. In many cases, subordinates will best conform to their moral obligations by deferring to their superiors, whose access to information, analysis, and legal counsel better positions them to make reliable proportionality judgments.

Of course, there are at least three situations in which subordinates should not defer to their superiors. First, subordinates may learn that the possible harm to civilians is much greater than their superiors anticipated. Presumably, a judgment based on false factual assumptions warrants little deference. Second, if the evidence available to subordinates suggests that an attack is *grossly* disproportionate then it is highly unlikely that additional evidence (available only to their superiors) would support the contrary conclusion. Finally, subordinates may have decisive reason to doubt that their superiors are applying the correct proportionality standard in good faith. In the first situation, subordinates should not carry out an attack unless they relay the new information to their superiors and then receive orders to proceed with the attack.²⁸ In the latter two situations, subordinates should not carry out an attack unless new information from their superiors convinces the subordinates themselves of its proportionality. As we shall see in Chapter 11, soldiers have no legal obligation, and no legal right, to follow orders that are manifestly illegal.

For example, suppose that a drone pilot is ordered to launch a missile at an unnamed, low-level insurgent away from an active battlefield. Before launching the missile, the pilot sees a dozen civilians pass into the expected blast radius of the missile. Presumably, the potential harm to civilians is much greater than the pilot's superiors anticipated. Moreover, the attack now appears so grossly disproportionate that it is highly unlikely that the pilot's superiors have specific information regarding the targeted individual that would support a judgment of proportionality. There is no reason to believe that the targeted individual will otherwise kill substantially more than a dozen soldiers or civilians, as would be required to render the attack proportionate. In this case, the pilot legally must and morally should (in the evidence-relative sense) suspend the attack.

Targeted Killings Revisited

²⁸ Compare Quéguiner, at 805 (“[I]f, before launching a first salvo against a bridge, a tank driver notices that a crowd of fleeing civilians have taken refuge under the targeted bridge, the driver cannot assume that the planners have correctly considered the principle of proportionality and continue his mission in wilful blindness and impunity. He must, at the very least, suspend his attack in order to allow the civilians to evacuate, or to request that his orders be confirmed in the light of these new circumstances.”)

Thus far, I have illustrated my account of *jus in bello* proportionality mostly using skeletal cases involving international armed conflict. While the U.S.-Iraq war shows that state-to-state conflict is hardly a thing of the past, it is worth pausing to apply my account to non-international armed conflicts between states and non-state actors.

As is well known, the United States has pursued a strategy of targeted killing of individual adversaries using unmanned aerial vehicles or “drones”. One of the supposed virtues of drones is that they can hover over a target for an extended period and strike when the fewest number of civilians are in harm’s way. Indeed, the overall percentage of civilians killed by drone strikes has fallen from almost 50% under the Bush Administration to under 20% under the Obama Administration, and the percentage has fallen still further in 2012 and 2013.²⁹ However, the Obama Administration has also expanded the drone campaign from targeting mostly high-level terrorists and insurgents to targeting primarily mid-level and low-level fighters. As a result, the number of drone strikes increased dramatically in 2010, and the annual number of civilian deaths remained very high. These changes in the design and execution of the drone program make it both more difficult and more important to evaluate its morality and legality.

To evaluate the proportionality of a particular drone strike, we must compare the losses the strike will inflict on civilians to the losses the strike will prevent from befalling soldiers or civilians over the remainder of the conflict. It therefore seems that drone strikes against low-level insurgents and terrorists that unintentionally kill civilians will seldom prove proportionate. Many low-level insurgents will never kill anyone, and many are so easily replaceable that killing them will not prevent substantially greater future harm. As we saw in Chapter 3, such insurgents may be liable to be killed to prevent them from harming soldiers or civilians, or causally contributing to others harming soldiers or civilians, even if they will be replaced if killed. In contrast, attacks that inflict losses on civilians can be justified only if the attacks also prevent substantially greater harm to others. It follows that if killing a low-level insurgent will not prevent substantially greater harm to others then it would be disproportionate to harm any civilians in the process.³⁰

In contrast, a high-level insurgent commander might make a necessary contribution to the deaths of many soldiers or civilians over the remainder of

²⁹ See Bureau of Investigative Journalism, *Covert Drone War*, <http://www.thebureauinvestigates.com/category/projects/drones/>.

³⁰ For a similar conclusion reached on strategic grounds, see DAVID PETRAEUS, U.S. DEP’T OF DEF., *THE U.S. ARMY AND MARINE CORPS COUNTERINSURGENCY FIELD MANUAL*, para. 7-32 (U.S. Army Field Manual No. 3-24, 2006) (‘In COIN environments, the number of civilian lives lost and property destroyed needs to be measured against how much harm the targeted insurgent could do if allowed to escape. If the target in question is relatively inconsequential, then proportionality requires combatants to forego severe action, or seek non-combatative means of engagement.’)

a conflict. High-level insurgents may be difficult to quickly replace, or may be replaced by others with less skill or charisma. In principle, it could be proportionate to kill some civilians as a side-effect of killing such a high-level insurgent. Importantly, such civilian losses would be unnecessary in the sense described in Chapter 7 if these losses could be avoided (for example by a ground operation) without placing soldiers or other civilians at comparable or greater risk. Moreover, even if a particular drone strike is proportionate in the *jus in bello* sense, the overall strategy of using drones to achieve broader war aims may be disproportionate in the *jus ad bellum* sense.

Conclusion

I have argued that an attack that inflicts harm on civilians is *jus in bello* proportionate only if it prevents substantially greater future harm to the attacking force or its civilians over the remainder of the conflict. This account of *jus in bello* proportionality does not compare incommensurable and imprecisely comparable values, only immediate losses to civilians and future losses to civilians and to attacking forces. In addition, this account applies symmetrically to all parties to an armed conflict, independently of the *jus ad bellum* morality and legality of their use of military force. Attacks that are disproportionate under this account are morally impermissible when carried out by just combatants, and disproportionate attacks carried out by unjust combatants are morally worse than proportionate attacks carried out by unjust combatants. It follows that both just and unjust combatants have decisive moral reasons to avoid attacks that are disproportionate under this account, and the law would guide soldiers well by prohibiting such attacks.