



The limited (but relevant) role of the doctrine of the double effect in the Just War Theory

Eduardo Rivera-López

To cite this article: Eduardo Rivera-López (2017) The limited (but relevant) role of the doctrine of the double effect in the Just War Theory, *Ethics & Global Politics*, 10:1, 117-139, DOI: [10.1080/16544951.2017.1396181](https://doi.org/10.1080/16544951.2017.1396181)

To link to this article: <http://dx.doi.org/10.1080/16544951.2017.1396181>



© 2017 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 14 Nov 2017.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

The limited (but relevant) role of the doctrine of the double effect in the Just War Theory

Eduardo Rivera-López

CONICET, Universidad Torcuato Di Tella, Buenos Aires, Argentina

ABSTRACT

My focus in this paper is the use of the doctrine of double effect (DDE) in the context of Just War Theory. Different versions of DDE have different degrees of plausibility. Two distinctions are crucial. First, I distinguish between epistemically idealized and non-idealized scenarios. Second, I distinguish between versions of DDE that make a *ceteris paribus* comparison between intentional and non-intentional outcomes, from versions that either make a comparison that is not *ceteris paribus* or are non-comparative. After undertaking these classificatory tasks, I defend the following claims: First, in an idealized world, *ceteris paribus* versions of DDE are plausible. Second, we cannot transfer such plausibility to stronger (non-*ceteris paribus*) versions of DDE. Finally, in a non-idealized world, DDE is plausible. The argument for this last claim combines the reasons for defending *ceteris paribus* versions of DDE in idealized circumstances with an argument about how to proceed when we face hard choices under uncertainty.

ARTICLE HISTORY

Received 8 May 2017
Accepted 5 October 2017

KEYWORDS

Doctrine of double effect;
just war theory; nonideal
ethics

1. Introduction

The doctrine of the double effect (DDE) remains, after centuries, a matter of controversy. In its long history, DDE has been used for different purposes in different discussions, ranging from self-defense to Just War Theory to bioethics. DDE has been formulated in different ways and with different scopes. My focus in this paper is the use of DDE in the context of Just War Theory; more concretely, the context of deciding to perform military actions that involve the use of lethal means affecting innocent persons. The conclusions of this paper may be, *mutatis mutandis*, useful for some other issues, although they may well not be useful in many cases.

Although the discussions around DDE are full of subtleties, I think that, at least on some dimensions, there is less than sufficient precision about DDE's content. DDE cannot be flatly endorsed or rejected. Different versions of DDE can have different degrees of plausibility. Two distinctions are crucial. First, we must distinguish between epistemically idealized and non-idealized (realistic) scenarios. Second, we must distinguish versions of DDE that make a *ceteris paribus* comparison between (roughly speaking) intentional and non-intentional outcomes, from versions that either make a comparison that is not *ceteris paribus* or are non-comparative. These clarifying and classificatory tasks are undertaken in sections 2 and 3.1. Next, I defend the following claims (I present these claims in a preliminary manner, acknowledging that they are still obscure): First, in an idealized world, *ceteris paribus* versions of DDE might

will be plausible (section 3.2). Second, we cannot transfer such plausibility to stronger versions of DDE (non-*ceteris paribus* and non-comparative versions) (section 3.3). Finally, in a non-idealized world, a strong version of DDE is plausible, but for different reasons. The (admittedly exploratory) argument for this last claim combines the reasons for defending *ceteris paribus* versions of DDE in idealized circumstances with an argument about how to proceed when we face hard choices under ineradicable uncertainty (section 4).

2. Basic assumptions and distinctions

As I said, I will discuss DDE in the context of Just War Theory. There are many controversial points in this theory, some of which might be relevant for the assessment of DDE. I will try, however, to skip over many of these by introducing several restrictions and assumptions at the outset of my argument. Let me mention the most important now.

- (1) I assume that radical pacifism is false. Some wars are morally justified (typically, a defensive war), and so are some actions in war.
- (2) I will constrain my discussion to actions performed in a just war. There is a huge controversy about whether actions in war can be morally justifiable in an unjust war, and, in general, about the independence of the *ad bellum* and *in bello* criteria of legitimacy.¹ I will not enter into this discussion. My assumption is that, *at least* in a just war, some actions in war are morally justified.
- (3) Any just war involves (at least a very high likelihood of) killing innocent persons. While there is enormous controversy about who should count as innocent, especially among civilians, I will take no position on this issue.² We can safely assume that in any just war there will be deaths among persons who are completely and undoubtedly innocent (children, civilians without any political commitment, etc.).
- (4) From points 1 and 3 it follows that there must be some killings of innocent persons (in a just war) that are not morally wrong. The discussion, then, is about the conditions under which killing innocent persons can be morally justified in a just war.

As is well known, DDE draws a distinction between two kinds of killing of innocents in war: intentional and non-intentional (merely foreseen). In order to illustrate this difference, I will make use of the following (modified) version of the traditional example that compares two kinds of bombings.³ Imagine you are a pilot (or a military commander). You are fighting a just

¹For the revisionist view, see McMahan (2009).

²For an outstanding discussion on this issue, see Fabre (2009).

³I modify the traditional examples comparing a 'terror bombing' (bombing civilians in order to demoralize the enemy and thereby obtaining a military advantage) and 'tactical bombing' (bombing a military facility thereby obtaining a military advantage, but killing civilians as a side effect). These traditional examples can be objected to for not fulfilling the *ceteris paribus* clause: something else, beyond intention, is different. In the first case, the military advantage is obtained indirectly (via demoralization), whereas in the second one it is obtained directly (via destroying a functioning military facility). Since it seems that the only way that killing civilians directly can contribute to obtaining a crucial military advantage is via demoralization of the enemy's leaders (or soldiers), the only way of making the comparison between the two bombings relevantly similar is to imagine that killing civilians indirectly (as a side effect) also contributes to the military advantage via demoralization of the enemy. This is why the example includes the destruction of a 'wonder weapon', which will demoralize the enemy. I might be objected that, if the wonder weapon is really harmless, the facility is not a legitimate military target (and the attack an act of vandalism). However, the distinction between military and non-military targets is not relevant here, because I am arguing at the level of the deep morality of war. I assume that attacking a harmless military facility can, under certain (very stringent) conditions, be morally permissible.

war. Victory is a morally valuable goal. Here are two possible actions you may consider performing (it may or may not be that both are available):

Foreseen Killing (FK): You bomb a secret military facility, where your enemy is developing what they believe to be a ‘wonder weapon.’ In fact, you know the weapon will be harmless. But, if you destroy the weapon, your enemy will be demoralized and your country will obtain a crucial military advantage. Unfortunately, some innocent persons are nearby and you will inevitably kill them if you bomb the facility.

Intentional Killing (IK): You bomb some innocent persons living in your enemy’s country. If you bomb and kill those persons, your enemy will be demoralized and, therefore, your country will obtain a crucial military advantage.

The difference between IK and FK stressed by DDE is that in IK the killing of the innocents is intentional, whereas in FK it is not. Of course, the concept of intentionality involved here is controversial. According to DDE, by ‘intentional’ killing we should understand that killing those persons is part of a plan of action in which their death is a means to obtain a goal. The action of killing is non-intentional when the plan of action does not include the death of those persons as a means. Their death is only a foreseeable side effect. The problem with this definition of ‘intentional’ and ‘non-intentional’ is that it already assumes the truth of DDE (it assumes that the morally relevant distinction is between being a means (or a goal) and being a side effect). It is extremely difficult to offer a definition that does not beg the question in favor or against DDE. My proposal is to assume, for the time being, the definition of intentionality just given (the one that would be accepted by defenders of DDE), which is useful to explain the different versions of DDE I want to distinguish. Later on, in [section 3.2](#), I will return to the concept of intentionality in the context of defending one version of DDE, and will offer a different, more plausible, understanding of the concept.

In order to assess the comparative and non-comparative features of IK and FK, we have to add two important considerations. The first is that there are many other respects, some of which are morally relevant, in which IK and FK might differ or be equal. The following is a list of some kinds of facts that might or might not be equal in both cases. I will call them ‘Relevant Facts’:⁴

- (1) Facts concerning the seriousness of the harm to be inflicted. These may include the number of fatalities produced by the bombing, number of injured persons, seriousness of the personal injuries, kind of victims (civilians, supporters, opponents, young people, old people, etc.), material damages, and other consequences in terms of the well-being of all involved persons (The moral costs).
- (2) Facts about the benefit to be obtained, for example, the number of persons to be saved from being killed, injured, or otherwise affected in their well-being if the bombing is carried out and the military advantage achieved (The moral benefits).
- (3) Facts concerning the availability of other, less harmful, courses of action to achieve the same valuable goal (Relevant Alternatives).
- (4) Facts concerning the probability of attaining the valuable goal (the crucial military advantage) (Probability of Success).

⁴Moreover, I will assume that, beyond facts about the agent’s intentions, these are the *only* morally relevant facts.

The precise content of the Relevant Facts and the agent's knowledge of them are crucial to determining the relative merits of FK and IK in a given situation.

A second important point concerns the degree of 'ideality' (or realism) we are going to assume in order to assess the relative or absolute merits of IK and FK. There are two kinds of scenarios I want to consider: ideal and nonideal. For my purposes, the distinction will be epistemic: an ideal circumstance (that I will call 'Ideal World') is one in which the agent knows the Relevant Facts with certainty.⁵ The 'Real World' is one in which decisions are made under realistic epistemic conditions of risk or uncertainty.

3. DDE in the Ideal World

In this section, I will distinguish between five versions of DDE in the Ideal World (3.1). For that purpose, I will use the FK/IK examples as a model of the distinction highlighted by DDE defenders. Although the classification may have some interest in itself, it will also be useful to distinguish two groups: those in which the contrast between FK and IK is *ceteris paribus*, and those in which this is not the case. After explaining the main features of the five versions, I will argue that we have (in the Ideal World) good reasons to accept *ceteris paribus* versions (3.2) and to reject non-*ceteris paribus* versions of DDE (3.3).

3.1 Five versions of DDE in the Ideal World

Suppose that we compare FK and IK, assuming a strict *ceteris paribus* clause: everything, except the facts explicitly mentioned in the example, is equal in FK and IK. The only relevant difference between FK and IK is that, in IK, the killing of the innocent persons is performed as a means to obtain the valuable military goal, whereas, in FK, the killing of innocent persons is only a foreseeable side effect of the means to the valuable military goal. Suppose further that, in this particular situation, FK and IK are the only available ways to obtain the crucial military advantage in the just war. A first formulation of DDE applied to war actions like FK and IK would be this:

1. *Comparative DDE (or, for short, C): assuming a strict ceteris paribus clause, and the absence of any other alternative to win the just war, FK is morally preferable to (or better than) IK.*

Regardless of the argument we might use in favor of C, it must be noted that C does not imply that FK is permissible and/or that IK is not. C is compatible with claiming that neither IK nor FK is ever permissible. They might both be always wrong, IK being worse than FK. C is also compatible with claiming that both actions (sometimes or always) are permissible, FK being preferable to IK. All this formulation of DDE says is that, if the agent faces the choice between FK and IK, and all other things are equal, it is morally better to choose FK. From the deontic point of view (that is, from the point of view of what is permissible, required, or

⁵The contrast between the Ideal and the Real Worlds should not be confused with the distinction between the deep morality of war and the conventions of war (McMahan 2004, 730). The whole discussion in this paper concerns the deep morality of decisions (in war). The point of section 4 and the discussion of DDE in the Real World is not how the conventions (or legal rules) should be like in an imperfect world, but how our moral decisions should be like in an imperfect world.

impermissible), the only plausible implications are that, *if* FK is not permissible, then IK cannot be permissible, and that *if* IK is permissible, FK must also be permissible.

The claim that FK is ‘morally better than’ or ‘preferable to’ IK is vague. One way to clarify these expressions is in terms of moral reasons. Saying that FK is morally better than IK (or that IK is worse than FK) could mean that there are some additional moral reasons against performing IK or some additional moral reason in favor of performing FK. For example, a defender of C might say that the fact that the killing is intentional in IK is a reason against performing IK that is absent in FK. Again, this does not mean that FK is permissible, because the Relevant Facts (which are assumed to be equal in both alternatives) might imply that both FK and IK are impermissible (that one has to accept defeat). Or it might be that both are permissible and FK is supererogatory.⁶

Although C might be considered the weakest (and therefore, least objectionable) version of DDE, it is not. Note that C claims that FK is morally preferable to IK, other things being equal, *regardless of how these other things are*. But how other things are, that is, how the Relevant Facts are, is not irrelevant and might affect the ranking, even if they are equal in both alternatives. For example, imagine an FK/IK pair of alternatives, in which the number of victims is equally disproportionate (the number of victims is the same in both alternatives, and is completely disproportionate in comparison to the moral benefits to be obtained). We might well think that, in this case (but not when the number of victims is more proportionate) IK and FK are morally indifferent. I am not affirming that this is the case. I just say that the Relevant Facts of the situation might well impinge upon the ranking of FK and IK, even if these facts are identical in both actions (and, therefore, the *ceteris paribus* clause is fulfilled).⁷

A weaker (and perhaps more plausible) version of DDE would constrain the Relevant Facts. One way to do so is to stipulate that some requirements or restrictions apply to them. For that purpose, it seems natural to think of the three conditions that are usually mentioned in the context of the *jus in bello* requirements of a just war: proportionality (facts 1 and 2), necessity (fact 3) and reasonable hope of success (fact 4). These three requirements have been extensively discussed and are beset by enormous complexities.⁸ I will not enter into them, because they are independent of the plausibility of DDE. For the sake of simplicity, I will assume that we have a plausible account of these requirements, in order to identify specific problems of DDE.

In the Ideal World, necessity and reasonable hope of success are superfluous requirements (not so in the Real World, as we shall see). Concerning necessity, I have stipulated that IK and FK are the only available alternatives for achieving the valuable goal; there is no less harmful alternative way to achieve the crucial military advantage. Hope of success is also irrelevant here, because I have stipulated that the

⁶For a purely comparative version of DDE, see Hills (2003, 134). Quinn’s version of DDE is also comparative and *ceteris paribus* (see Quinn 1989, 344).

⁷This is why, in my view, Nathanson’s argument against DDE fails (Nathanson 2010). He argues that DDE implies that, had the intention of the perpetrators of the September 11 attack to the Twin Towers not be killing people but just destroying the buildings, their action would have been less wrong. This is, as Nathanson rightly points out, obviously false. However, the example is one in which not only there is no recognizable valuable goal pursued by the terrorists, but also (and for the same reason) there is no sense of proportionality in the action. The phenomenon I am calling attention to here is a case of what Kamm calls ‘contextual interaction’ (see Kamm 1993, 53 ff).

⁸For a very nuanced analysis of proportionality in war, see (Hurka 2005). For simplicity, many of my examples will be focused only on innocent victims, that is, on what Jeff McMahan calls ‘wide proportionality’ (see McMahan 2009, 20–21). On necessity in war, see Lazar (2012).

agent knows with certainty that both IK and FK are efficient ways to achieve the valuable goal.⁹ Therefore, the only relevant requirement in the Ideal World is proportionality: whether IK and FK cause a proportionate harm in comparison to the benefit to be achieved.¹⁰ For simplicity, I will call the three mentioned requirements (proportionality, necessity, and reasonable prospects of success) ‘Relevant Requirements’ (although, as I said, in the Ideal World only proportionality is relevant). The Relevant Requirements apply to the Relevant Facts, establishing how those facts must be in order for IK or FK to be acceptable candidates to be morally justified actions (they establish a necessary condition, not a sufficient one).

With these considerations in mind, we can formulate a weaker, and more plausible, version of DDE for FK/IK kinds of circumstances:

2. *Constrained Comparative DDE (or CC): assuming a strict ceteris paribus clause, the absence of any other alternative to win the just war, and that the Relevant Requirements hold, FK is morally preferable to (or better than) IK.*

The same remarks I have made about C apply here. CC is also compatible with saying that both alternatives are permissible or with saying that neither is (for example, it is compatible with radical pacifism). Moreover, here too, we can conceptualize the expression ‘morally preferable’ as meaning that there is a reason against IK that is absent in FK (the fact that in IK the killing is intentional). Among the relevant versions of the DDE, CC is the weakest and safest one.

We can make CC stronger in the following way:

3. *Deontic Constrained Comparative DDE (DCC): assuming a strict ceteris paribus clause, the absence of any other alternative to win the just war, and that the Relevant Requirements hold, FK is morally permissible and IK is not.*¹¹

DCC is still a comparative doctrine. With some trivial assumptions, we might plausibly infer from DCC that, in a scenario where FK alone is available (and the Relevant Requirements hold), FK would be permissible. However, we cannot infer that, in a scenario where IK alone is available, IK is impermissible, because it might be that the prohibition on IK (when FK is available) is connected to FK’s availability.

We can conceptualize the difference between CC and DCC by resorting again to the language of reasons. We can say that DCC assumes that the reasons against doing IK (or in favor of doing FK) are sufficiently strong to make IK impermissible (and FK permissible). But they only make IK impermissible, *provided* that FK is available. As said, nothing follows

⁹Moreover, some authors plausibly think this requirement can be subsumed by proportionality (Hurka 2005, 37).

¹⁰In the Just War Doctrine (as well as in the rules of the humanitarian international law), proportionality only applies to collateral harms to civilians. This assumes some version of DDE. Since I do not want to beg the question in favor of DDE, I will consider that both IK and FK can be proportionate (or disproportionate), in the sense of producing (or not) an adequate relationship between humanitarian harms and benefits.

¹¹For a formulation of DDE that is similar to DCC see (McIntyre 2001, 219) (‘It is sometimes permissible to bring about a harm as a merely foreseen side effect of an action aimed at some good, even though it would have been impermissible to bring about the same harm as a means to that end.’).

about whether IK still be impermissible were it to be the only way to achieve the valuable goal of victory.

We can make DCC still stronger by dropping the *ceteris paribus* condition:

4. *Unrestricted Deontic Constrained Comparative DDE (UDCC): assuming the absence of any other alternative to win the just war, and that the Relevant Requirements hold, FK is morally permissible and IK is not, even if everything else is not equal.*

I will discuss this version more carefully in the next section. The important point for now is that, although UDCC is not a *ceteris paribus* version of DDE, it still is, in one sense, comparative. It says that, given the availability of IK and FK, and *provided* the Relevant Requirements hold, you are allowed to choose FK (and not allowed to choose IK), regardless of the differences about the Relevant Facts. Such a view would still be compatible with holding that, were IK the only way to win the just war, it would be, in some occasions, permissible.

Given the characterization of the Relevant Requirements I offered above, UDCC is particularly relevant in the Ideal World when the degree of satisfaction of the first requirement, proportionality, is different in FK and IK. Even if two actions are proportionate, one of them can be closer to the upper limit of proportionality than the other. I will return to this point in the next section.

Finally, we can conceive of a last version of DDE:

5. *Deontic Constrained Absolute DDE (DCA): assuming that the Relevant Requirements hold, FK is morally permissible and IK is not, even if everything else is not equal, or only one of the alternatives is available.*

DCA is not comparative, because it does not assume that both alternatives are available. Even if the only way to obtain victory is IK, we are forbidden to proceed. We must accept defeat by the unjust enemy over violating this principle. This version of the DDE is very similar to the traditional understanding of the doctrine: there is a moral constraint against doing intentional harm, even if doing so is necessary to achieving a valuable goal. However, there are cases in which doing harm is permissible. These cases are those in which the harm is not intentional (is a side effect), and the Relevant Requirements are fulfilled.¹² In terms of reasons, we can say that, according to DCA, there is a *decisive* reason against killing innocent persons intentionally, but the reasons against killing can be defeated in a just war, when the Relevant Requirements hold and the killing is not intentional in the relevant sense.

3.2 In favor of *ceteris paribus* versions of DDE

Let me now argue in favor of *ceteris paribus* versions of DDE (CC and DCC). Remember CC: faced with the situation in which FK and IK are the only available options, all other things being equal, and provided that the Relevant Requirements obtain, there is a moral reason to prefer FK over IK. Once this has been established (3.2.1), I will offer an additional argument in favor of DCC: in the same circumstances and under the same conditions, it is permissible to perform FK and not permissible to

¹²Scanlon discusses (and rejects) this version. See (Scanlon 2008, 1; 20–32).

perform IK (3.2.2). In both cases, it is worth emphasizing that the versions of DDE to be defended are rather weak: they compare two kinds of actions (FK and IK) under very specific and demanding conditions: (i) that the agent is absolutely certain about all the consequences of both courses of action, (ii) that both actions have exactly the same consequences in terms of harms and benefits (which means that the Relevant Requirement are equally satisfied in both cases), and (iii) that FK and IK are both available and are the only ways to obtain a valuable goal (victory or crucial military advantage in a just war). In a sense, what the arguments in favor of CC will try to prove is just that there is a relevant moral reason to ‘break the tie’ between FK and IK in favor of FK.

3.2.1 *Defending CC*

One traditional argument favoring DDE is that IK implies a kind of involvement of the agent that is qualitatively different from, and more serious than, the involvement implied in FK. In IK, the agent involves herself in harming innocent victims in a deeper way. She is ‘*aiming at evil*,’ in the sense that bombing innocent people becomes, in a way, a goal (or, as Alison Hills says, a plan-relative end). She will adjust her behavior in order to achieve such a goal (for example, if people change their location, the pilot will have to change the location of the bombing in order to reach them).¹³ In a way, she hopes that the civilians do not escape from the bombing. On the other hand, when innocent deaths are a side effect, no such adjustment is necessary: if the potential victims relocate outside the reach of the bombs, so much the better. The agent, in this case, is not *aiming at killing* them. Even if there is zero likelihood that the civilians can escape the bombing, she hopes they would. The difference in personal involvement in evil has, according to this view, some relevant consequences. For example, by killing people as a side effect, the agent wants their death in a different manner from when it is a mere instrument to her goal. Therefore, she can apologize to the victims (or their relatives) in a different, more plausible, way. She can face the fact of their innocence with more dignity, by explaining that their death was not an essential part of her plan.

Against this argument, there have been numerous objections. Without trying to recall all of them now, it seems that one key problem is that the distinction between intentional and foreseen effects heavily depends on how we describe the agent’s action. We can always find a description of the events according to which the relevant differences between FK and IK disappear. For example, we can say that, in FK, the pilot’s intention is to produce the destruction of the military facility *and* the death of the civilians. Since, under the circumstances, both occur together, she cannot intend one without intending the other. The two events form a ‘package.’ If this is right, the pilot is ‘*aiming at evil*’ in FK as well. For the same reason, the kind of apology cannot be substantially different. The victims (or their relatives) will not find any comfort in the fact that their death was not part of the plan, because, in one relevant sense, it was: it was something that the pilot accepted as a reasonable moral cost, much in the same way as the pilot in IK did. Alternatively, we might say that the deaths are side effects in *both* cases. In the case of IK, the agent can say that the death of the civilians is a side effect as well: what she *intends* is merely for the leaders of the enemy’s country to *believe* that the

¹³See Nagel (2001), Hills (2003), Smith (2007, 352).

civilians were killed, but should the civilians manage to miraculously survive without their leaders knowing it, so much the better! This is all that is required to produce their demoralization and surrender. Their actual death is merely a side effect of causing that belief.¹⁴ As is well known, there are two responses to this kind of objection. The first is to claim that the death of the victims in cases of direct intention is ‘too close’ to the intended effect (the civilians appearing to be dead). The second response is to widen the concept of intentional action and speak not of killing the victim but of harmfully involving her. In what Quinn calls ‘harmfully direct agency’ I involve the potential victim intentionally, as a means to an end. For example, I involve the civilians intentionally, even if I do not intend their death.¹⁵

In what follows, I try to find an alternative (and complementary) argument, specifically focused on supporting CC. The challenge is, firstly, to single out the fundamental difference between IK and FK and, secondly, to explain why it bears (at least some) moral weight. Let us consider each question in turn.

The general idea is that intention comes in degrees. Therefore, the effects of our actions can be more or less intentional. This depends on the more or less central role they play in our plan of actions. In the case of FK and IK, the relevant effect is the civilians’ death, and we have to determine how intentional that effect is within the pilot’s (or the commander’s) plan. The purpose is not to establish something like an objective degree of intentionality according to some scale, but to be able to compare pairs of cases, like FK and IK, in order to say whether the civilians’ death is, in one case, ‘more intentional’ than in the other, and why. In order to do so, I draw to the following counterfactual test.¹⁶

Let us call ‘X’ the event we are most concerned with in situations like FK and IK: the death (or harm) of innocent civilians. Our purpose is to determine how central X is within our plan of action. The first thing to do is to imagine possible worlds in which we (i) perform the relevant action (FK or IK), (ii) the goal is achieved, and (iii) X does not obtain. In the case of FK, this means imagining possible worlds in which we bomb the military facility, demoralize the enemy’s leader, and thereby achieve the goal of victory (or military advantage), without killing the innocent civilians. In the case of IK, that means imagining possible worlds in which we bomb civilians, demoralize the enemy’s leader, and thereby obtain the crucial military advantage, without killing the innocent civilians. Let us call those possible worlds that fulfill conditions (i) to (iii) ‘happy’ possible worlds (‘happy’ because no person dies).

Among all the happy possible worlds, we should compare each case (FK and IK) with the happy possible world (or set of possible worlds) that is *closest* to the actual world. Let us call those (closest) happy possible worlds FK* and IK*.

If we try this exercise, we will note that FK* will normally be closer to FK than IK* is to IK. In the case of FK, the closest happy possible world FK* is one in which everything is equal, except that civilians are not in the wrong place and time (perhaps they are just one meter away from it). In the case of IK, the closest happy possible world IK* is more distant. We have to imagine a world in which I bomb the city and yet they do

¹⁴See Bennett (1995, 210–211).

¹⁵Quinn (1989, 343–344). Tadros (2015) is a sophisticated development of Quinn’s view. For a useful discussion of both alternatives, see Fitzpatrick (2006, 587–8).

¹⁶I take the test from Kamm (2000, 26) and (2004: 654), although she does not endorse it (she thinks it is unable to draw the relevant distinction of her theory of triple effect).

not die, but the leaders believe that they did. The hypothetical causal chain leading from bombing civilians (who do not die) to the demoralization (due to belief in their death) is far more complex and involves more changes to the actual world than we need in order to obtain FK* from FK. One reason why IK* is more distant from IK than is FK* from FK is that, in IK, the civilians' death is part of the causal explanation of the pursued goal (military victory): their deaths explain the leaders' demoralization and their surrender. If we want to explain the leaders' demoralization and surrender by our bombing without the civilians' death, a different causal chain must be imagined, with (many more) different facts involved. All this is not necessary with FK.

The distance between the happy IK* and FK* worlds and the actual IK and FK worlds is also indicative of the number of happy possible worlds that are associated with each alternative. This is important because it reflects in how many counterfactual scenarios I would change or adjust my behavior. If the event X (the death of the civilians) is more important to my plan, I will change or adjust my behavior in all alternative scenarios (scenarios in which X does not occur) except in a few specific ones (those where the very unlikely 'happy' situation arises, in which nobody dies and victory is achieved). On the other hand, if X is less important for my plan, I will not change my action in most alternative scenarios, because there are many happy possible worlds in which I bomb the military facility and no civilian dies. In all those alternative scenarios I will bomb the facility nonetheless.

Let us assume my counterfactual test is effective to comparing the degree of importance of an effect within the structure of our plan of action. The second, crucial, question is whether this difference between FK and IK is sufficient to allow us to say that I have some moral reason to choose FK and not IK. Remember our assumption here is that the pilot (or the country she is working for) is allowed to defend her country and that FK and IK are the only ways. Therefore, she *has* to choose among them. Is FK morally better?

My answer is 'Yes': she has more moral reasons to do FK. The death of human beings being intrinsically bad, bringing about such an intrinsically bad event as a more important component of our plan of action is worse than bringing it about as a less important component, *ceteris paribus*.

In order to see why this central point inclines the moral balance in favor of FK, let us reconsider the two classical considerations mentioned above, about the kind of involvement of the agent and the kind of apology the agent can provide in either case.

Imagine again the FK and IK pilots expressing their intentions. The FK pilot can credibly say 'I am aiming at destroying the military facility. Unfortunately, there are 100 civilians that will be killed by my bombing.' Instead, it would be extremely contrived for the IK pilot to say 'I am aiming at causing the enemy to believe that 100 civilians were killed. Unfortunately, there is no other way to achieving this other than actually killing them.' The same happens with the consideration about the apology that each offer to the victims or their families. In FK, the pilot can provide this apology: 'I am very sorry. The civilians' deaths were not my intention. Killing them was not in my plan. It is an unfortunate fact that they were close to the military factory. Had they not been there, I would have bombed the factory anyway.' The IK pilot should say something like this: 'I am very sorry. The civilians' death was not my intention. Killing them was not in my plan. It is an unfortunate fact that they were such that the bombs killed them instead of just making them appear to be dead. Had the bomb just produced the mere appearance of their death, I would have bombed them anyway.' This is obviously a much less

credible apology. It is less credible because the pilot should admit (and is not admitting) that the civilians' death was important for her plan, or at least more important than in the case of FK. In both cases (the expression of her intentions and the apology), the FK pilot is comparatively more credible, because the death of the civilians is less central to her plan. Therefore, the counterfactual ('happy') scenario, in which no civilian dies is more realistic. The counterfactual test is the tool that shows how central the harmful effect of our action is to our plan.

We have to remember that, in order to evaluate CC, the relevant question now is not whether FK or IK are permissible, but rather whether, with these (and only these) two actions available, it is *better* to choose FK rather than IK, or is it a matter of indifference. The question is whether, the options being exactly identical, except for the mentioned counterfactual element about the centrality of the death of civilians for the agent's plan, there is any reason to do FK that is morally relevant and is absent in the case of IK. I think there is one. The reason is slim, in a way, but sufficient. We do not want to kill. Therefore, we should arrange our conduct in ways in which killing is less relevant for our plan. Both in FK and IK, the pilot kills the civilians, but killing them is less central in FK than it is in IK. She is less involved, at least in the sense that, by opting for FK, she expresses more of an attitude of rejection or disapproval to the act of killing. All other things being equal, she has a moral reason to prefer that way of acting.

The fact that what is at issue in FK/IK cases is just a pairwise moral comparison between two options has implications for some arguments against DDE. Howard Nye (in Nye 2014) has recently proposed a complex argument, in which he tries to go from the famous Trolley case (in which we are allowed by DDE to turn the trolley from the main track with five persons to a side-track with only one person) to the (also famous) Transplant case (in which we are not allowed by DDE to kill one person and use his organs to save five persons) along several intermediate cases, claiming that there is no relevant moral difference between each case and the next. The conclusion is that there is no relevant moral difference between Trolley and Transplant and, therefore, that DDE should be rejected. I will not discuss each step of the argument (there are 11), but only two of them. The sixth example is equivalent to FK (Nye calls it 'The Big Bomb') and the ninth is equivalent to IK (called 'Standard Terror Bombing'). Nye tries to go from one to the other over two intermediate examples. I will consider one of them, called by Nye 'Lewissian Terror Bombing'.¹⁷ For the sake of clarity I will modify the example, making it more similar to IK and FK:

Lewissian Killing (LK):

You explode fragmentation bombs in the air over a number of innocent civilians. The explosion will cause the enemy's leaders, who will see the explosion from afar, to surrender, but it will also kill those innocent civilians.¹⁸

LK is similar to FK in that the pilot does not exactly seek the civilians' death. In a way, the death of the civilians can be considered a side effect (the pilot in LK only seeks that the enemy's leaders see the explosion and believe they are dead). On the other hand, it is similar to IK in that in both the civilians are more directly involved than in FK (or in Nye's 'The Big Bomb'). The crucial point, in my view, is that, if the pilot

¹⁷The example is attributed to David Lewis by Quinn (Quinn 1989, 343).

¹⁸I am paraphrasing Nye's version of the example (Nye 2014, 461).

should have to choose between FK and LK (always assuming ideal epistemic conditions and a strict *ceteris paribus* clause), the pilot has more moral reasons to choose FK, because the death of the civilians is less central to her plan in FK than in LK. By the same token, if a pilot should have to choose between LK and IK, the pilot should opt for LK. Although in both the death of the civilians is more directly involved in the pilot's plan (compared with FK), in IK the death of the civilians is the cause of the enemy's belief they are dead. In my previous terminology, FK* (the nearby happy possible world in which the valuable goal is achieved and nobody is harmed) is closer to FK than LK* is to LK, because 'removing' the innocent civilians in FK (that is, imagining a world without there being innocent civilians there) involves fewer changes than removing them in LK (removing them in LK implies not only imagining they are not there, but also that the enemy believes they are there).

Similarly, the fact that CC makes pairwise comparisons makes my argument compatible with Kai Draper's recent objections to DDE. Draper claims that DDE is unable to make 'strongly discriminating' distinctions between intentional and non-intentional harm. It allows only 'weakly discriminating' distinctions. But these are useless as a guide to making strategic decisions at war (Draper 2015, 123–124). According to Draper, strongly discriminating formulations of DDE hold that some kinds of actions are 'much' harder to justify than some other kinds. My view is that the difficulty in justifying action A over action B is a matter of degree and depends on the specific features. What matters here is that actions of killing in which the victims' death is more central to our plan are more difficult to justify, *ceteris paribus*, than actions of killing in which the victims' death is less central. This is enough to make CC plausible under ideal epistemic circumstances. What I try to show in section 4 is that this 'weakly' discriminating principle is not useless when we apply that principle to the real world, in which radical uncertainty pervades war actions.

3.2.2 Defending DCC

As we have seen, CC (that there are some reasons to prefer FK over IK) does not say anything about the permissibility of FK and/or IK. It might be that both actions are wrong or that both are right. Let me then argue for DCC: if CC is true and only FK and IK are available, FK is permissible and IK is not, provided that the Relevant Requirements are satisfied and all other things are equal.

We can first set aside the possibility of both actions being always wrong. They can be wrong in some cases but not always, since we have assumed from the start that radical pacifism is false. The alternative would be that both FK and IK are permissible, with FK preferable to IK. This position would imply claiming that FK is a kind of supererogatory action: we have a moral reason in favor of FK (or against IK) but, still, morality leaves us a certain freedom. However, I do not think this is acceptable. Those who defend the existence of supererogatory actions think that they are not obligatory because they are, to some extent, costly for the agent. They are too demanding for the agent required to perform them. They are 'beyond the call of duty.' This is also why the agent becomes laudable if she performs the supererogatory action. But if the cost of the action is exactly the same either way (which is ensured by the *ceteris paribus* clause), it would be very odd to think that we should not opt for the action which we have stronger reasons to perform. Imagine that I can

either donate 100 or 80 dollars to some valuable cause. Suppose that, if I donate 80, the remaining 20 would be lost anyway for some reason, so that the cost for me would be the same. In this case, it seems clear that it is not only the case that I have a reason to donate 100 and not 80. I *should* give 100. I am required to opt for the action which stronger reasons speak for.

In the case of DCC we have an analogous situation. First, we have some moral reasons to prefer FK over IK. Second, the cost of performing IK and FK is the same. Third, everything else is equal. It seems sensible to conclude that, having assumed CC, FK is not only the preferable alternative but the only permissible one.

3.3 Against non-ceteris paribus DDE versions in the Ideal World

Suppose you accept that *ceteris paribus* versions of DDE are plausible: given the choice between FK and IK, there is a reason to choose FK and there is no reason to choose IK.

Can the plausibility of *ceteris paribus* versions of DDE (especially of DCC) be extended to non-*ceteris paribus* and non-comparative versions (UDCC and DCA)? The substantial point is whether the moral *ceteris paribus* superiority of FK over IK is strong enough to overcome other factors that may appear when we drop the *ceteris paribus* condition, or when we consider FK and IK separately.

Let us consider UDCC (when only FK and IK are available, FK is always permissible and IK never is, even if everything else is not equal, provided that the Relevant Requirements hold). Imagine the following FK/IK situation. You have to choose between:

Facility Bombing 80: you bomb a military facility, where the enemy is developing what they falsely believe to be a wonder weapon. If you destroy that facility, your enemy will be demoralized and you will win the war. Unfortunately, 80 children of a nearby school will inevitably die if you bomb the facility.

School Bombing 40: you bomb a school, where 40 children are present. Since they are the children of the regime's leaders, if you bomb that school, those leaders will be demoralized and you will win the war.

Let us assume that 100 innocent fatal victims is the highest number that, in this particular war, is acceptable if we want to satisfy proportionality (one of the Relevant Requirements). Let us further assume that everything else is equal. According to UDCC, you are morally allowed to perform Facility Bombing 80 (and not School Bombing 40). Note that CC and DCC do not have this implication, since they compare IK and FK stipulating that all other things (besides the fact that the children die as a means in one case and as a side effect in the other) are equal, *including* the number of children. This example, on the contrary, assumes that all other things are equal, *except* the number of victims (and the fact that the children die as a means in one case and as a side effect in the other).

The crucial question is whether the argument called for in order to convince us that CC and DCC are true in the Ideal World is sufficient to convince us that UDCC is true and, therefore, that we have to choose Facility Bombing 80.

In Facility Bombing 80, there are 40 additional victims. We should somehow compare the weight of the death of these persons against the weight of the fact that the deaths of the 80 victims are not intentional in Facility Bombing 80 and the deaths of the 40 victims are intentional in School Bombing 40. How do we make such a comparison?

I have defended both CC and DCC as plausible. That means that, if you faced the choice between Facility Bombing 80 and, what we may call, School Bombing 80, in which everything is just like it is in School Bombing 40 except that the number of victims is 80 (that is, equal to Facility Bombing 80), you would be only permitted to perform Facility Bombing 80, and not School Bombing 80.

On the other hand, we can imagine a further case, Facility Bombing 40, in which everything is like it is in Facility Bombing 80, but the number of (unintentional) victims is 40. Suppose you have to choose between Facility Bombing 80 and Facility Bombing 40. In this case, Facility Bombing 40 is obviously preferable to Facility Bombing 80. Furthermore, I would say that *only* Facility Bombing 40 is morally permissible (Facility Bombing 80 would violate the necessity requirement). That means that, everything else being equal, numbers count: it is preferable to kill 40 persons to killing 80. Killing more persons is not irrelevant. It matters, and matters heavily, even if the number of victims is proportionate in both alternatives. If we accept this, we cannot say that the fact that 40 additional persons will die in Facility Bombing 80 (against School Bombing 40) is irrelevant, or that it provides us with no reason at all in favor of School Bombing 40.

The outcome thus far would therefore be the following. On the one hand, we have a moral reason to prefer Facility Bombing 80 over School Bombing 40 (the same reason that makes us prefer Facility Bombing 80 over School Bombing 80): that the harm, in Facility Bombing 80 is more intentional, more central to our plan of action. I will call this reason the ‘intentionality reason.’¹⁹ On the other hand, we have a moral reason to prefer School Bombing 40 to Facility Bombing 80: that the harm, in Facility Bombing 80, is greater in terms of the number of persons you will kill (I will call this reason ‘harm reason’).

It might seem that balancing these two kinds of reasons against one another is impossible because they are incommensurable: intentionality reasons are deontological, whereas harm reasons, based on the number of victims, are consequentialist. Still, I will argue that a comparison must be, at least to some extent, possible.

In order to balance the involved reasons properly, we have to keep seriously in mind that the choice between Facility and School bombings is made in the Ideal World. That means that you are absolutely certain about all the Relevant Facts, and you know that everything is equal, except the intentional structure and the number of victims. You know that, in both cases, the Relevant Requirements hold. And you know with certainty the effects of your actions. Under these highly idealized conditions, compare these two theoretical options:

- (1) Radical DDE: intentionality reasons prevail over every possible harm reason in the Ideal World. That means that, whatever the number of additional victims involved, we should always choose Facility Bombing over School Bombing (always provided that the Relevant Requirements hold).
- (2) Moderate DDE: in *ceteris paribus* comparisons (where the harm is the same), only intentionality reasons matter. In non-*ceteris paribus* comparisons (where the harm is different), harm reasons become relevant as well and can, up to a certain degree of harm, overcome intentionality reasons.

¹⁹ I am assuming that the reason to choose FK in *ceteris paribus* circumstances remains operative in non-*ceteris paribus* ones. This is plausible in my view, but not obvious. In any case, assuming this claim gives support to the versions of DDE I want to reject.

According to Moderate DDE, the harm reason to choose IK and not FK increases in strength (beginning with zero, when the harm is the same in both options) as the difference in harm increases between the options. The intentionality reason to choose FK and not IK, on the contrary, remains constant, whatever the harm implied in each alternative. If we think that harm reasons are relevant, there must be a point at which harm reasons prevail over intentionality reasons. I do not have a theory that establishes the amount of harm (i.e. the number of innocent victims) necessary to incline the decision in favor of IK. I can only offer the following thoughts. In the case of *ceteris paribus* versions of DDE (CC and DCC), the agent (the pilot) can offer a justification to the victims of FK: she has chosen the option that did not imply aiming at killing innocent victims as a relevant part of her plan of action (or better: she has chosen the option that implied killing the innocent victims as a less relevant part of her plan). Given the strict equality in harm, victims of FK would not have anything to complain about. On the other hand, in non-*ceteris paribus* comparisons, they do have something to say: ‘we were additional victims! You could have spared some deaths.’ On the other hand, the pilot opting for School Bombing 40 does have a good explanation for his intentional killing, which is not available in *ceteris paribus* pairs of examples: sparing a relevant number of innocent deaths.

How many additional victims are necessary to incline the balance in favor of IK is a very difficult question. It depends upon how strongly we think intentionality reasons should impact our moral deliberation. In concordance with the argument previously developed in favor of CC, I would suspect that just one additional victim is sufficient to override intentionality reasons. However, I do not want to press this here, because it is not strictly necessary to rejecting UDCC. Even if we think that intentionality reasons are more than tiebreakers, there must be *some* amount of harm able to defeat those reasons. Therefore, we cannot hold non-*ceteris paribus* versions of DDE without qualification.

What about DCA? According to this version of DDE, let us recall, IK is never permissible and FK, provided that the Relevant Requirements hold, always is (in the Ideal World). The implausibility of UDCC infects DCA. The argument is this. I have claimed that UDCC is implausible because, at least in some cases, harm reasons prevail over intentionality reasons. Suppose that my example is one of those cases: we have to choose School Bombing 40 over Facility Bombing 80 (in the Ideal World). Assuming this is the case, now imagine that we have to choose between School Bombing 40 and nothing. We know, by hypothesis, that, if Facility Bombing 80 had also been available, School Bombing 40 would be the right choice (and therefore morally permissible). It would be odd to claim that, just because Facility Bombing 80 is not available, School Bombing 40 becomes impermissible. It must be permissible.

That means, in clearer terms, that, in the Ideal World, if the Relevant Requirements are satisfied, some terrorist war acts are morally permissible, when (and only when) less harmful (or equally harmful) non-terrorist acts are not available.²⁰

The conclusion thus far is the following. *Ceteris paribus* versions of DDE (CC and DCC) are plausible in the Ideal World. We have intentionality reasons that, when

²⁰Can this conclusion be applied to classical examples involving killing and letting die? For example, if I claim that some terrorist attacks are permissible in the Ideal World, am I forced to say, for example, that a doctor is allowed (in the Ideal World) to kill an innocent bystander to transplant her organs to five dying patients? Not necessarily. There is at least one important difference. The cases of terrorist attacks that would be morally permissible are those for which a (real or counterfactual) alternative action of killing (an action that kills as a side effect) is accepted as morally permissible. In other words, School Bombing 40 is morally permissible because (among other things) Facility Bombing 80 would also be permissible (and this is not a disputed claim). In the transplant case, there is no parallel, undisputed, case.

everything else is equal, carry some moral weight, and incline the choice in favor of not inflicting lethal harm intentionally. On the other hand, the non-*ceteris paribus*, comparative, version of DDE (UDCC) is implausible in the Ideal World. When there is a difference in the amount of harm among the relevant alternatives, we have a strong reason to minimize harm, which can defeat intentionality reasons. Our reason to minimize harm may incline the choice in favor of IK, when the (proportionate) harm is lower (or at least significantly lower) than the harm to be inflicted by FK. Finally, and for the same reasons, the non-*ceteris paribus*, non-comparative, version of DDE (DCA) is also implausible: in the Ideal World, and provided that the Relevant Requirements strictly hold, IK may also be permissible when it is the only option to win the just war.

4. DDE in the Real World

4.1 From the Ideal to the Real World

Real human decisions are not made in the Ideal World. In my view, this fact has received less attention than deserved in moral philosophy, especially by non-consequentialist moral philosophers. This lack of attention is surprising if we note that the conclusions we may arrive at for the Ideal World are not obviously directly applicable in the Real World (our world) and may even be irrelevant or misleading.²¹ We need criteria to adapt ideal conclusions to real circumstances.

As we have seen in [section 2](#), the distinction between the Ideal and the Real World is epistemic: in the Ideal World we have assumed perfect knowledge about the consequences of our actions. In the Real World, this perfect epistemic capacity does not exist. The degree of reliability of our knowledge in the Real World depends on several factors: the time spent in deliberation, the capacity to obtain the relevant evidence, the capacity to evaluate the evidence impartially (without bias, external pressures, etc.), the magnitude and diversity of the foreseeable consequences, among others. In all these variables, decision-making in war ranks *very* low: decisions are taken with enormous time pressure, the capacity of impartial deliberation is dramatically impaired by the fact that these are decisions taken by one party against an enemy, and the potential magnitude and diversity of consequences are huge. Just taking into account these few factors, it seems quite obvious that the beliefs of decision-makers about long-term consequences of their actions in war are, at the moment of taking those decisions, extremely unreliable.²²

²¹For exceptions, see Hansson (2010, 586–592); Spector (2016).

²²Concerning the case of proportionality (one of the Relevant Requirements), Amichai Cohen seems to be pointing to this kind of problem when he says: 'The most obvious difficulty generated by the concept of proportionality relates to the feasibility of rational choice in battlefield situations. In essence, the concept of proportionality assumes that the attacker has an option. He can either take action, thereby bringing to bear his full military advantage, or he can relinquish a small part of his military advantage, in order to significantly limit the damage that he will cause to civilians. Confronted with that choice, proportionality requires the attacker to choose the latter course. In reality, however, military operations very rarely offer such a clear-cut choice. In most cases, the dangers and potential of the battlefield are unknown. The commanding officer simply cannot make an informed assessment of the harm to the civilian population vis-a-vis the gains from a specific military attack. Hence, the problem with proportionality is that it seems to place an unrealistic burden on the attacker to make a cost-benefit analysis in the midst of a battle.' (Cohen 2010, 12). In the same direction, this is how the so-called NCV (Non-Combatant Cutoff Value) is established, according to the former military lawyer (now teaching at Texas University), Geoffrey Corn: 'There's no benchmark. And there's nothing in the law or doctrine, like a table you would open up ... [and says] a tank is worth two civilian casualties, or a command post is worth five civilian casualties. [...] Most of the science behind NCV is left up to commanders' subjective analysis' (Watkins 2016). I thank an anonymous reviewer for calling my attention to Watkins' article.

With this preliminary picture of the dramatic epistemic shortcomings of decision-making in war, we might seriously ask: are our conclusions about DDE in the Ideal World translatable into the Real World? Are those conclusions relevant? If they are, to what extent? What are the adjustments or modifications necessary to adapt ideal moral criteria to nonideal circumstances? These are pressing and fundamental questions. I would like to offer a preliminary exploration of them, by way of suggesting the claim that, in view of the severe epistemic limitations decision-makers have in the Real World about the morally relevant consequences of belligerent actions, DDE might have a more relevant role in the Real World than we might expect.

4.2 *Uncertainty and consequentialism*

The problem of uncertainty of long-term consequences of actions is not limited to the ethics of war. It has been presented as a general ‘epistemic argument’ against consequentialism.²³ According to this argument, since long-term consequences of our actions are unpredictable, we can never know whether our actions have optimal consequences. For all we know, any action can have, in the long run, the best or the worst morally relevant overall consequences. We would have no clue, then, which action is in fact right or wrong and, therefore, consequentialism would not serve as a decision procedure. It would be no basis for moral guidance.

Assume for a moment that the epistemic argument is convincing. We might try to apply this general argument to conduct in war. Since our harm reasons (our reasons not to cause humanitarian harm to others in war) are time-neutral, we should take into account the consequences in their entirety, not just the short-term, foreseeable, effects. But, since the proportion of foreseeable effects is insignificant in comparison to the totality of consequences, we would have no clue about how to act on the basis of harm reasons.

If we assume this epistemic argument and my conclusion in [section 3.2](#) about the relevance of intentionality reasons in *ceteris paribus*, ideal, situations, we might construct an argument to show that, under realistic conditions of decision-making in war, we should be sensible to intentionality reasons. The idea would be that, harm reasons being useless (given their inability to guide our action), we are left only with intentionality reasons. DDE would, in that way, be vindicated in the Real World.

I do not think this argument is right, stated this way, among other reasons, because it would prove too much: not just that harm reasons are inert in the case of war actions, but that no decision at all should rest on consequentialist reasons. Still, I think a more nuanced and sophisticated version of the argument may well be plausible, or at least worth exploring.

One way to reject the epistemic problem is to claim that, long-term consequences being radically uncertain, we should simply dismiss them. We should assume that positive and negative unknown effects of our actions cancel each other out, leaving us with those short-term effects that we are able to foresee. The claim that unpredictable effects will cancel each other out rests on the principle of indifference: when we are genuinely uncertain about two outcomes, we should attribute equal probabilities to them. However, there are well-established objections to this principle, and defenders of

²³See Lenman (2000, 342–343). Lenman takes Shelly Kagan’s formulation of the problem.

consequentialism usually accept these objections and try to find other ways to encounter the epistemic problem.²⁴

One of these ways is to deny that the epistemic shortcomings of moral agents are as radical as the objection assumes. It might be that, as Burch-Brown argues, certain kinds of actions are statistically correlated to bad consequences and others to good consequences. For example, we might have reasons to believe that, statistically, allowing the indiscriminate killing of innocent people on the street is more likely to produce bad consequences than prohibiting it (Burch-Brown 2014, 116). A crucial point of this argument is to move from the level of individual actions to the level of patterns of behavior, strategies, or policies. We cannot know whether an isolated action will have long-term optimal consequences, but we can know that actions that fall under a certain pattern will more probably be more (or less) harmful than actions belonging to a different pattern. And this may be sufficient to guide our action in consequentialist terms. Still, there may be patterns that we can rank and patterns we are unable to rank, depending on the context in which decisions are made. In what follows, I will consider the specific context of war.

4.3 *Radical uncertainty among strategies*

At the beginning of this section (4) I mentioned some features of decision-making in war: lack of time to deliberate, strong bias in favor of one's own cause (lack of impartiality), enormous magnitude of impact of decisions on key aspects of the life of many people, extreme difficulty in gathering relevant evidence about the situation and the consequences of possible ways of action, among others. These features make decision-making in war especially vulnerable to epistemic shortcomings and, therefore, to the necessity of finding general guides, or rules, that we have reasons to believe would rank better as to the fulfillment of the Relevant Requirements: proportionality, necessity, and prospect of success.

There are different patterns or strategies we might think best in terms of the Relevant Requirements. Following my previous distinction between harm reasons and intentionality reasons, we can say that these (consequentialist) patterns will only be backed by harm reasons, reasons to minimize the harmful effects of war in the long run. Intentionality reasons would play no role. However, we might now ask whether there are (harm) reasons to prefer a pattern of behavior in war that includes intentional harming (e.g. targeting civilians) over a pattern that excludes such kind of harm (restricting war actions to targeting only military objectives). In other words, we should compare the following two possible patterns or rules:

- (A) In order to maximize the fulfillment of the Relevant Requirements, military decision-makers will be allowed to attack civilians, as well as military targets.
- (B) In order to maximize the fulfillment of the Relevant Requirements, military decision-makers will only be allowed to attack military targets.

²⁴See Lenman (2000, 351–359); Mason (2004, 318); Burch-Brown (2014). Very roughly, the problem of the principle of indifference is that from the lack of knowledge about probabilities we cannot infer knowledge about equal probabilities, as much as we cannot infer knowledge about any other distribution of probability.

The distinction between (A) and (B) imperfectly tracks the distinction between harms that are more intentional (actions in which the immediate resulting harm is more central to the plan of the decision-maker) and harms that are less intentional (actions in which the immediate resulting harm is less central). The distinction might be traced differently. The important question is whether we have any (harm) reason to think that strategy (A) will have better humanitarian consequences than strategy (B) in the long run. I think we are epistemically unable to answer this question affirmatively. In fact, some have argued that there is a consequentialist argument in favor (B) (not attacking civilians intentionally). The idea would be that, if every belligerent adopts, as a strategy or rule of behavior ‘never attack civilians intentionally; do so only as a side effect,’ the humanitarian long-term effects of war will be, other things being equal, better. This sounds reasonable and I do not want to reject it out of hand. Still, it is an extremely bold empirical claim, which I do not know how to support with evidence.²⁵ In any case, if the claim is true, it would give reasons in favor of a consequentialist rule akin to the DDE. But my argument for DDE does not rest on this empirical claim. At most, I assume a much weaker empirical claim, namely, that there is no strong evidence that (A) is a better strategy, morally speaking. Given the extreme uncertainty about the long-term harmful (and beneficial) consequences of war actions, I think we should conclude that strategies (A) and (B) are, at least in the vast majority of cases, indifferent from the point of view of harm reasons.

4.4 *Intentionality reasons becoming salient*

Compared with the empirical knowledge about long-term consequences of actions, intentions are relatively transparent, at least for the agent’s own intentions. In most cases, the pilot (or military commander) can confidently determine, for example, when the death of innocent civilians in a war has been planned as a mean to achieve some end and when it was a (foreseeable) side effect. There may be gray cases, or cases in which the agent has both intentions. However, the intention of killing innocent civilians directly is a relatively clear and discrete one, and can be easily identified by the agent. There is of course a legitimate philosophical discussion about the kind of access we have to our own intentions (beliefs and desires).²⁶ Still, it seems safe to assume that our access to our own intentions is by far more reliable than our access to the whole number of long-term consequences of a lethal attack in a war.

If we accept this, a military commander would be, in the Real World, in most cases, in the following situation:

- (1) He is unable to make a case-by-case estimation of the long-term humanitarian consequences of each available action in the different war scenarios he must face. He is not able to rank the alternative, eligible actions along the Relevant Requirements.

²⁵Howard Nye quotes some research showing that terrorism is not efficient to achieve the goals terrorists pursue (see Nye 2014, note 1). But not every act of intentionally killing innocent civilians qualifies as a terrorist act.

²⁶For a discussion on (and defense of) transparency of intentions, see Byrne (2011), esp. section 6.

- (2) He is able to make a consequentialist (harm-based) ranking of different patterns or rules, according to general statistical estimations about normal consequences of following those different patterns.
- (3) He has no (harm) reason to believe that systematically including intentional harm into the menu of (*prima facie*) permissible actions will have better humanitarian consequences than excluding intentional harm.
- (4) At least in clear cases, he is able to distinguish those possible actions in which the harm to innocent persons is intentional and those in which the harm is not intentional, at least in the approximate way of distinguishing military targets from killing civilians directly.

This being so, my claim is that the military commander (or, in general, a decision-maker in war) has moral reasons to adopt the rule of avoiding intentional harm. The version of DDE that is closest to this picture is the last, strongest one, DCA: within the range of eligible actions, a military commander should always opt for military targets, and never attack innocent civilians intentionally.

I hasten to grant that, in the Real World, we cannot exclude the possibility of exceptional cases. For example, suppose we are nearly certain that an intentional harm is much more proportionate than a non-intentional one and the difference in the amount of harm were huge, or that we are nearly certain that no non-intentional harm is possible under the circumstances. We would then probably say that we are morally permitted to harm intentionally (provided that the Relevant Requirements hold). But in most real war scenarios, uncertainty is endemic and, therefore, intentional harm to innocent civilians will be harder to justify. From the individual point of view of a military commander, the conclusion thus far is that, given the extreme unreliability of the available predictions about harms and benefits, he should follow a rule that allows intentional lethal harm only in exceptional cases, when the disparity of harm is huge and can be confirmed just by appealing to rough estimations of consequences.²⁷

4.5 *The residual requirement to minimize harm*

The following objection should be considered. If epistemic unreliability obliterates any difference in the dimension of harm (so that only intentionality reasons remain salient), then we have no reason to choose, among non-intentional harmful actions, those that produce less harm (according to the available evidence). This is extremely implausible. On the other hand, if we think (as I do) that there is a requirement of minimizing harm (harm reasons *are* operative in the Real World as well), then why not adopt this precept from the start? It seems that, if the estimations necessary to minimize harm are possible (even if they are imperfect), and therefore, judgments about magnitudes of harm (and benefits) are meaningful, then harm reasons would make sense in the Real World and would defeat intentionality reasons, exactly as they do in the Ideal World, according to my argument in [section 3.2](#).

²⁷This may have implications at the institutional (or conventional) level, to the effect of supporting rules prohibiting intentional killing innocent persons in war. My point, however, is not about the conventions that can be morally justified to have. It is a deep morality point about the moral duties of agents in war (see footnote 4 above).

The crux of the problem seems to be why, on the one hand, the epistemic unreliability of fine-grained, long-term estimations of harm causes harm reasons (in the Real World) to be defeated by intentionality reasons, but, on the other hand, once intentionality reasons have done their work (by discarding intentional harms), harm reasons start to be operative again (by requiring to minimize non-intentional harm).

My answer is that, when we try to establish the morally relevant harms and benefits of an action of war, the immediate number of victims is not the only relevant factor. We should take into account all immediate and more distant effects. This calculus is what I think is impossible to carry out in a minimally reliable way. We do not have any reason to think that systematically allowing both intentional and non-intentional harm in the pursuit of minimizing humanitarian harm (or of optimizing the satisfaction of the Relevant Requirements) will be more efficient for that purpose than allowing only non-intentional harm. Therefore, we have intentionality reasons to adopt a strategy excluding intentional harm. This is compatible with claiming that, once this has been settled, we also have moral reasons to adopt other strategies, which, according to rough general estimations, would most probably satisfy the Relevant Requirements.

5. Conclusion

As Thomas Nagel rightly pointed out, our moral life involves both deontological, agent-relative reasons and consequentialist, impersonal reasons.²⁸ There is, in my view, no general algorithm to give the exact weight to each kind of reason in every relevant situation. We have to argue, in each kind of situation, for the best balance.

An important part of the literature on moral theory tries to reach normative answers by arguing on the basis of highly idealized examples. The discussion on DDE is no exception. I find those examples useful, insofar we are fully aware of the (unrealistic) conditions under which those examples should be thought to obtain, and of the (limited) scope of the conclusions we are able to reach with their help. The outline of the argument in favor of DDE, at the idealized level (what I have called Ideal World), has been the following:

- (1) We assume that, provided that the Relevant Requirements are satisfied, some killings of innocent persons are morally allowed.
- (2) We assume that, provided that the Relevant Requirements are satisfied, FK is morally allowed.
- (3) The question is, then, whether, with FK and IK available, and all else being equal, IK would also morally be permissible.
- (4) With all else equal, the intentionality reason becomes salient: we have reasons to opt for FK. IK is, therefore, impermissible.

Then I considered the case in which (in the Ideal World) not everything else is equal, for example, because (and only because) in FK more innocent victims will be killed. In this case, my point was that reasons about minimizing the death of innocent victims prevail over intentionality reasons. In non-*ceteris paribus* (ideal) situations, DDE becomes morally inert.

²⁸See Nagel (1986, 164–188).

Finally, I addressed the problem of how all this impinges upon our real, nonideal world. In a fully exploratory spirit, I have tried to motivate two basic points. First, at least in the case of war, we should be extremely skeptical about the ability of combatants to predict and calculate all the morally relevant consequences of their actions. They are only able to endorse strategies or rules of action according to rough and imperfect statistical estimations. And there is no reason to believe that strategies including intentional harm to innocent civilians are better than strategies excluding that kind of harm, with regards to the fulfillment of the Relevant Requirements. This being so, and allowing the possibility of exceptions, the strongest version of DDE becomes operative in the Real World. If a belligerent decision-maker cannot be sure about the moral cost of his lethal actions, he can at least be sure to be acting in a more moral way.

Acknowledgements

For helpful comments on earlier versions of this paper, I am grateful to Yuval Abrams, Alejandro Chehtman, Marcelo Ferrante, Josep Maciá, Thomas Schmidt, Ezequiel Zerbudis, to the audiences at the Universidad Pompeu Fabra, Universidad Torcuato Di Tella, and Sociedad Argentina de Análisis Filosófico, and to two anonymous reviewers.

Disclosure statement

No potential conflict of interest was reported by the author.

References

- Bennett, J. 1995. *The Act Itself*. Oxford: Oxford University Press.
- Burch-Brown, J. M. 2014. "Clues For Consequentialists." *Utilitas* 26 (1): 105–119. doi:10.1017/S0953820813000289.
- Byrne, A. 2011. "Transparency, Belief, Intention." *Aristotelian Society Supplementary Volume* 85 (1): 201–221. doi:10.1111/j.1467-8349.2011.00203.x.
- Cohen, A. 2010. *Proportionality in Modern Asymmetrical Wars*. Jerusalem: Jerusalem Center for Public Affairs.
- Draper, K. 2015. *War And Individual Rights: The Foundations Of Just War Theory*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199388899.001.0001.
- Fabre, C. 2009. "Guns, Food, and Liability to Attack in War." *Ethics* 120 (1): 36–63. doi:10.1086/649218.
- Hansson, S. O. 2010. "The Harmful Influence of Decision Theory on Ethics." *Ethical Theory and Moral Practice* 13 (5): 585–593. doi:10.1007/s10677-010-9232-0.
- Hills, A. 2003. "Defending Double Effect." *Philosophical Studies* 116: 133–152. doi:10.1023/B:PHIL.0000005744.80837.65.
- Hurka, T. 2005. "Proportionality in the Morality of War." *Philosophy & Public Affairs* 33 (1): 34–66. doi:10.1111/j.1088-4963.2005.00024.x.
- Kamm, F. 1993. *Morality, Mortality. Volume II*. Oxford: Oxford University Press.
- Kamm, F. 2000. "The Doctrine of Triple Effect and Why a Rational Agent Need Not Intend the Means to His End. Part I." *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 74, pp. 21–39.
- Lazar, S. 2012. "Necessity in Self Defense and War." *Philosophy & Public Affairs* 40 (1): 3–44. doi:10.1111/j.1088-4963.2012.01214.x.
- McIntyre, A. 2001. "Doing Away with Double Effect." *Ethics* 111: 219–255. doi:10.1086/233472.

- McMahan, J. 2004. "The Ethics of Killing in War." *Ethics* 114: 693–733. doi:10.1086/422400.
- McMahan, J. 2009. *Killing in War*. Oxford: Oxford University Press.
- Nagel, T. 1986. *The View from Nowhere*. Oxford: Oxford University Press.
- Nagel, T. 2001. "Agent-Relative Morality." In *The Doctrine of Double Effect*, edited by P. A. Woodward. Notre Dame: University Of Notre Dame Press.
- Nathanson, S. 2010. *Terrorism and the Ethics of War*. Cambridge: Cambridge University Press.
- Nye, H. 2014. "On The Equivalence Of Trolleys And Transplants: The Lack Of Intrinsic Difference Between 'Collateral Damage' And Intended Harm." *Utilitas* 26 (4): 432–479. doi:10.1017/S0953820814000181.
- Quinn, W. 1989. "Actions, Intentions, and Consequences: The Doctrine of Double Effect." *Philosophy & Public Affairs* 18 (4): 334–351.
- Scanlon, T. 2008. *Moral Dimensions. Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
- Smith, I. 2007. "A New Defense of Quinn's Principle of Double Effect." *Journal of Social Philosophy* 38: 349–364. doi:10.1111/josp.2007.38.issue-2.
- Spector, H. 2016. "Decisional Nonconsequentialism and the Risk Sensitivity of Obligations." *Social Philosophy & Policy* 32 (2): 91–128. doi:10.1017/S0265052516000121.
- Tadros, V. 2015. "Wrongful Intentions without Closeness." *Philosophy & Public Affairs* 43 (1): 52–74. doi:10.1111/papa.12043.
- Watkins, A. 2016. "The Numbers Game." *BuzzFeed News*. https://www.buzzfeed.com/alimwatkins/syria-civilian-casualties-policy?utm_term=.nh3VyPG16#.lj3YeOQBj.