Weighing unjust lives

Forcehimes, Andrew T.

2017

Forcehimes, A. T. (2017). Weighing unjust lives. In O. Jens David, M. Larry, & F. Claire (Eds.), *Weighing lives in war* (pp. 284-297). doi:10.1093/oso/9780198796176.003.0013

<https://hdl.handle.net/10356/145501>

<https://doi.org/10.1093/oso/9780198796176.003.0013>

© 2017 Andrew Forcehimes. Published by Oxford University Press. All rights reserved. This book is made available with permission of Andrew Forcehimes.

Downloaded on 26 Apr 2023 23:13:15 SGT

Weighing Lives in War

Edited by
JENS DAVID OHLIN
LARRY MAY
and
CLAIRE FINKELSTEIN

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© the several contributors 2017

The moral rights of the authors have been asserted

First Edition published in 2017

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Crown copyright material is reproduced under Class Licence
Number C01P0000148 with the permission of OPSI
and the Queen's Printer for Scotland

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2017945335

ISBN 978-0-19-879618-3 (pbk.)

ISBN 978-0-19-879617-6 (hbk.)

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Weighing Unjust Lives

Andrew T Forcehimes

1. Introduction

Are the lives of those fighting on the unjust side of a war worth less than the lives of those fighting on the just side? To get a grip on this question, consider

Single Death: Country U unjustly starts a war with Country J. U and J stockpile weapons and rally troops. However, U, given J's defensive capabilities, spends much of the war posturing. When U and J meet on the battlefield, few shots are fired, mostly for strategic purposes. Whenever U advances on J's territory, J always forces a hasty retreat. The war drags on. Unable to make substantial progress, U signs a peace treaty with J. In tallying up the costs, it is discovered that, shockingly, only a single combatant was killed, hit by a stray mortar.

In *Single Death*, we can assume that the combatant is deprived of many years of happy life. Such deprivations are bad. But are they worse if we learn that this combatant fought for J rather than U? All else equal, it is tempting to answer yes. If this stray mortar had to hit someone, and we were in a position to choose, the choice seems obvious. It would be better if it hit the unjust.

There is an attractive rationale for this verdict: Things are intrinsically better when people get what they deserve. Being on the just side, all other things equal, the combatant deserved to bear the cost of dying less than her unjust counterpart. Call this view, which treats the goodness of a life as the product of one's desert-adjusted welfare, the *desert-adjusted account*.

There is much to be said in favor of adjusting for desert. People differ in their levels of desert, and this should make a difference. If I am vicious and you are not, it is better that your welfare increase over mine.¹ Similarly, if one of us has to suffer, my viciousness counts strongly in favor of it being me over you. These claims seem correct. And the desert-adjusted account provides a tidy explanation, hence its popularity.²

¹ Following S Kagan, *The Geometry of Desert* (Oxford University Press 2012), when talking about desert I will use the language of virtue and vice. This is simply for convenience. Nothing substantive is intended. Virtue is simply the placeholder for whatever it is that makes one more deserving.

² See F Feldman, *Confrontations With the Reaper: A Philosophical Study of the Nature and Value of Death* (Oxford University Press 1992); J McMahan, *The Ethics of Killing: Problems at the Margins of Life* (Oxford University Press 2002); Kagan (n 1). Kant also held this view, it seems. He writes, "What now is the *summum bonum*? It is the unification of the greatest happiness with the greatest degree of capacity to be worthy of this happiness. If there is to be a highest good, then happiness and the worthiness thereof must be combined ... If we conduct ourselves in such a way that, if everyone else so conducted themselves, the greatest happiness would arise; then we have so conducted ourselves as to be worthy of happiness." Quoted from P Guyer, *Kant on Freedom, Law, and Happiness* (Cambridge University Press 2000) 93.

Adjusting for desert provides a powerful justification for weighing the lives of the unjust differently from those of the just. If correct, the desert-adjusted account has significant implications for just war theory.³ In war, we are forced to weigh lives. Such calculations require an account of what justifies assessing one death as worse than another. The desert-adjusted account tells us that it is worse for the deserving to be deprived of future welfare than the undeserving. Since an unjust combatant has a lesser moral status, her death is less bad than the death of a just combatant (assuming both are deprived of equal amounts of welfare).⁴

Given that weighing lives pervades nearly every aspect of war, it is important to know if the desert-adjusted account is true. Despite its appeal, I believe that the account is mistaken. The argument proceeds in four sections. Section 2 fleshes out the desert-adjusted account and explains why many find it attractive. Sections 3 and 4 are critical: Section 3 highlights the troubling implications that adjusting for desert has for *ad bellum* and *in bello* proportionality calculations; Section 4 turns to the accounts implications for cases in which the welfare of noncombatants cannot be influenced, but their level of viciousness can be. The implausibility of these implications calls into question the core idea of the desert-adjusted account: namely, that there is some level of welfare that each person deserves, and things would go best if everyone were at these levels. Section 5 concludes.

2. The Desert-Adjusted Account

To understand the desert-adjusted account of weighing lives, we need to understand the badness of death. On the standard account—*deprivationism*—the badness of death is a comparative bad.⁵ Imagine two possible life trajectories for the combatant in *Single Death*. On Path A, the combatant is hit by a stray mortar and dies. On Path B, the mortar does not stray and she lives. After her time in the army, she goes on to have a reasonably pleasant life, dying of natural causes in old age. Path A is bad because of how it compares with Path B. Path A is cut short. The combatant suffers a loss. She is deprived of what Path B affords.⁶ The badness of death is traceable to this difference.

³ A now-standard line in just war theory is that the unjust have made themselves liable in ways that just combatants or non-combatants have not: see J McMahan, "The Basis of Moral Liability to Defensive Killing" (2005) 15 *Philosophical Issues* 386. But notice, the desert-adjusted account tells us something different. It makes an axiological claim: Unjust lives count for less. Killing the unjust, that is, impacts the total amount of intrinsic good to a lesser extent than killing the just. Perhaps adjusting for desert offers a deeper explanation than rights forfeiture. If consequentialists hope to get an extensional fit with standard just war theory, the desert-adjusted account offers an attractive option.

⁴ Unless noted, the equal deprivation of welfare should be assumed. It should also be assumed that the death takes away a good. We can leave aside cases where, because her life would have only negative welfare, the person's death is a welcome relief. For brevity, I will not always make these assumptions explicit.

⁵ For the classic defense of deprivationism, see T Nagel, "Death" (1970) 4 *Noûs* 73. For a recent defense with important clarifications, see J Broome, "The Badness of Death and the Goodness of Life," in B Bradley, F Feldman, and J Johansson (eds), *The Oxford Handbook of Philosophy of Death* (Oxford University Press 2013).

⁶ A number of complications concerning counterfactuals—e.g. problems of overdetermination—are here left aside. I am inclined to a "nearest possible worlds" approach. For a defense of this approach, see F Feldman, "Some Puzzles About the Evil of Death" (1991) 100 *Philosophical Rev* 205. For critique, see McMahan (n 2) 119.

But what, in particular, does death deprive us of? Is there more than one factor?⁷ If there are multiple factors, how do they interact?

One obvious factor is welfare. If welfare is the only factor lost in death, then we can say that the badness of a person's death is solely a function of the welfare her life would have otherwise contained. Call this the *straight-welfare approach*. On this approach, if we want to know the badness of *Single Death*, we compare the welfare contained in Path A (where the combatant dies) with her welfare in Path B (where she lives). The difference in total welfare between Path A and B tells us the badness of her death.

Cashing out the badness of death solely in terms of welfare fits well with many of our considered judgments. For example, it seems intuitively worse for a twenty-five-year-old to die than an eighty-year-old. Why? Because, assuming a plausible welfare per-year distribution, the twenty-five-year old loses more welfare than the eighty-year-old. That seems like the right verdict.

Still, problems loom for the straight-welfare account. Consider

Factory: There has been an explosion at the local factory, and Andy and Becky have both sustained life-threatening injuries. Without immediate medical treatment, they will die. We are EMTs on the scene. We have medicine that can save the life of either Andy or Becky, but not both. We know, however, that Andy culpably caused the explosion—after being passed over for a promotion, he intentionally left the gas on. All other things are equal.⁸

In stipulating that all else is equal, we stipulate that the welfare contained in Andy's and Becky's lives, were they not to die, is exactly the same. This stipulation, combined with the view that welfare is the only relevant factor in determining the badness of death, tells us that Andy's and Becky's impending deaths would be equally bad. If we knew that helping Andy or Becky would have no indirect effects—for example, incentivizing immoral behavior—we could decide whom to save by flipping a coin. Surely that would be a mistake. Even if we could keep Andy's behavior entirely secret, it seems we should help Becky. But, if we think it would be better to save Becky than Andy, then we must think that the badness of death does not consist in the deprivation of welfare alone.

Enter desert-adjusted weighing. If we want to hold that it is better to save Becky, we need to identify a worse-making feature that sets her death apart from Andy's. The main difference, of course, is that Andy is less virtuous than Becky. In light of Becky's greater virtue, she deserves to be better off. Our medicine will prevent a loss of welfare for one of these people. The world, we might plausibly say, is a better place if Becky is given the welfare, as opposed to culpable Andy. According to the desert-adjusted account, then, the badness of death is a function of the amount of welfare one receives and the amount of welfare one deserves. If we adopt this account, we can maintain, all

⁷ McMahan (n 2) 183, for example, lists six factors. I will not discuss all six. But I will discuss his fifth, "Were the individual's previous actions or character such as to make him deserving of the goods he lost?" See *ibid.* That McMahan endorses this factor commits him to the desert-adjusted account.

⁸ Modified from Kagan (n 1) 23.

other things equal,⁹ that we should help Becky, because Andy's culpability makes his death less bad than Becky's.

We should not, however, be misled by the specifics of *Factory*. The desert-adjusted account is not concerned with doling out specific rewards (or punishments) for specific virtues (or vices). Instead, the account is interested in the overall—bottom-line—amount of welfare one receives. The desert-adjusted account, in other words, is a *whole-life approach*.¹⁰ As Kagan clarifies, “[W]e look at lives as a whole, to see what one deserves (overall), and whether one has received it (overall).”¹¹

Visualizing the differences between a straight-welfare and a desert-adjusted approach will be helpful. It will also allow other important details of the desert-adjusted account to emerge. In *Factory*, we can bestow a fixed amount of welfare on either Andy or Becky. We want to know if things go better or worse if we give the medicine to Becky over Andy. If we let the X-axis indicate welfare levels (with increasing welfare from west to east)¹² and let the Y-axis indicate intrinsic goodness of a given state of affairs (with goodness increasing from south to north), then the straight-welfare account's answer to this question can be represented graphically as is shown in Figure 12.1.¹³

Fig. 12.1 Straight-Welfare

⁹ Notice that the “all other things equal” stipulation now tells us that Andy and Becky stand to lose the same welfare *and* that their desert levels, were they not to die, hold steady.

¹⁰ On this whole-life approach, it matters if, say, Becky experienced a wonderful childhood with unusually high levels of welfare, while Andy had a particularly nasty childhood with very low levels of welfare. If Becky and Andy's desert levels were close enough and the gap in prior welfare was wide enough, then our verdict might be reversed. For now, we can put this complication aside and assume rough equality between their prior-to-*Factory* welfare.

¹¹ Kagan (n 1) 11. See also Feldman (n 2) 182.

¹² Points to the east of the origin are lives worth living. Points to the west of the origin are lives not worth living. I am concerned, at this point, only with the former. I draw the graphs accordingly.

¹³ In representing the competing approaches graphically, I follow standard practice in the desert literature. See, for example, F. Feldman, “Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice” (1995) 55 *Phil and Phenomenological Research* 567; T. Hurka, “The Common Structure of Virtue and Desert” (2001) 112 *Ethics* 6; and most impressively Kagan (n 1). Many of the graphs below are modified from Kagan (n 1) ch 3.

As Figure 12.1 shows, on the straight-welfare account, Andy and Becky's lines overlap. Any deprivation in welfare is also a deprivation in intrinsic goodness. It is good that we give the medicine to Andy or Becky: It doesn't matter whom. Either distribution is equally good. We should be indifferent between the two.

The desert-adjusted account tells us otherwise. As Kagan puts it,

[F]or each person there is an absolute level that the person deserves to be at. This is what the person deserves absolutely. If people have what they deserve, this is good ... If people have less than they deserve, then this is less good, or perhaps even bad ... [If] someone has more than they deserve, this is less good, or perhaps even bad.¹⁴

We can capture the core tenets of the desert-adjusted account with three principles.

Peaks: The line for each level of virtue is mountain-shaped, with a highest point. This peak is what each person, in terms of overall welfare, deserves.

Shift: If Becky's virtue level is greater than Andy's, then the peak welfare level for Becky is greater than the peak for Andy.

Bell Motion: If Becky's virtue level is greater than Andy's, then the mountain for Becky rotates counterclockwise.¹⁵

Andy has a certain virtue level. So does Becky. But, given Andy's culpability, Becky's virtue level is higher. Becky is more deserving. Each of these virtue levels has a corresponding peak. The peak for Andy is what he deserves given his virtue level. Becky also has a peak. Like Andy, her peak indicates the level of welfare she deserves. *Peaks* thus tells us that welfare contributes to intrinsic goodness up to a point.¹⁶ This point is set by one's level of desert. After this point, welfare's contribution to goodness drops off.¹⁷ So, if Andy and Becky are at their respective peaks, any change in welfare is a change for the worse.¹⁸ However, given *Shift*, Becky's peak corresponds to a greater level of welfare than Andy's. In slogan form, "the more deserving deserve more."¹⁹

Peaks and *Shift*, taken together, tell us that different amounts of intrinsic goodness correspond to different levels of welfare for Andy and Becky.²⁰ They thus set the desert-adjusted account apart from the straight-welfare account. Yet recall our stipulation that Andy and Becky would be deprived of the same amount of welfare. Suppose this amount is exactly what it would take to get Andy to his peak. Would the desert-adjusted account then tell us to be indifferent? To avoid answering yes, we need *Bell Motion*.

¹⁴ S Kagan, "Equality and Desert" in L P Pojman and O McLeod (eds), *What Do We Deserve* (Oxford University Press 1999) 300.

¹⁵ I have modified this way of capturing the desert-adjusted account from B Skow's review of S Kagan, "The Geometry of Desert" (2014) 124 *Ethics* 417, 418.

¹⁶ The straight-welfare approach, by contrast, tells us that it is always good the more welfare one receives.

¹⁷ Kagan (n 1) 82.

¹⁸ Ibid.

¹⁹ Ibid.

²⁰ *Peaks*, it is worth noting, does not favor retributivism—though, obviously, it is compatible. Figure 12.2 is neutral between retributivists and non-retributivists. A retributivist holds that it can be good for the vicious to suffer. That is, retributivists affirm, while non-retributivists deny, the existence of negative peaks—that is, peaks that fall to the west of the origin on the X-axis.

Bell Motion captures an intuitive thought: The more virtuous you are, the more important it is that you receive the welfare you deserve. Consequently, if Becky is more virtuous than Andy, it is worse for Becky to be below her peak by a certain amount of welfare than it is for Andy to be below his peak by that same amount.²¹ *Bell Motion* ensures that the virtuous convert welfare into goodness more efficiently than the vicious.

We can now represent the desert-adjusted account graphically as follows.

Fig. 12.2 Desert-Adjusted Account

As Figure 12.2 illustrates, if both Andy and Becky possess welfare below their peaks, distributing additional welfare to Becky increases intrinsic goodness more than distributing welfare to Andy. Becky has a steeper western slope. *Bell Motion* also tells us that if Becky gets more than she deserves, it is less bad than if Andy gets more than he deserves. If, for example, the welfare benefits of continued existence overshoot Becky's peak, we should still give her the medicine—for, given her more gradual eastern slope, it is less bad for Becky to get more than she deserves than Andy.²² It is, in a word, better to shortchange the vicious and overcompensate the virtuous.²³

Peaks, *Shift*, and *Bell Motion* form the core of the desert-adjusted account. The desert-adjusted account delivers the intuitively correct verdict in *Factory* and *Single Death*. No doubt, the account is enticing. When death deprives someone of future welfare, better that it fall on the wicked. That is hard to deny. Nevertheless, we should be wary. By turning to the context of war, the problems with the desert-adjusted account become vivid. In the next section, I focus on *Peaks*.

²¹ Kagan (n 1) 103.

²² *Bell Motion*, of course, allows for the possibility that it is better to help the more vicious. We have already seen one possibility, given the whole-life approach. Another possibility is that we have unequal amounts of welfare to distribute—e.g. we can give Becky only a few years of happy life but Andy many years of happy life. I am, for now, assuming these complications away.

²³ Kagan (n 1) 100.

3. The Problem with *Peaks*

Peaks, recall, tells us that there is some level of overall welfare each person deserves. If one is at her peak, things are going as well as they could. Any increase or decrease in welfare makes things worse. *Peaks* lies at the heart of the desert-adjusted account. But consider what *Peaks* implies for wide-proportionality calculations—calculations involving those who have not made themselves liable to be killed.²⁴

Wide-proportionality calculations are intimately connected with lesser-evil justification. When making such calculations we compare the good achieved by acting in a certain way with the bad produced by acting in this way. The good effect must be sufficiently good to compensate for the bad effect. Wide proportionality tells us, *ad bellum*, that the bads inflicted on the non-liable must not be out of proportion with the goods produced by going to war. *In bello*, it tells us that a specific act of war, say, the bombing of a munitions factory, is impermissible if the bad caused by the bombing—for example, the collateral killing of the non-liable—is out of proportion with the goods achieved by bombing.

With this rough sketch,²⁵ we can draw out the implications of the desert-adjusted account. Consider a variant of the *in bello* case just described

Eastern Civilians: Country U is unjustly fighting a war with Country J. J knows that U is housing a stockpile of weapons in a factory within bombing distance. These weapons will bring the war to a favorable end sooner, but not by much. J thinks the bombing is worth pursuing, but wants more information on the collateral damage. Reliable reports reveal that 1,000 of U's non-combatants will be killed if J bombs the factory, but the lives of all of these non-combatants are, given their desert levels and welfare levels, east of their peaks. And, despite the war, the future for these non-combatants looks bright.

Wide proportionality tells us to refrain from acts in cases where marginal gain comes with tremendous loss. But, if the desert-adjusted account is true, *Eastern Civilians* is not such a case. The noncombatants' welfare levels are east of their peaks. Their welfare levels outstrip their virtue levels. So, assuming that their further lives will have on-balance positive welfare, things only get worse if U's noncombatants live. Their deaths, according to the desert-adjusted account, would make things go better than their continued existence. Wide proportionality would thus not rule out J's bombing of the munitions factory. We could not say, if the desert-adjusted account were true, that J's bombing is impermissible because the badness caused by the bombing is out of proportion with the goodness achieved. Why? Because the deaths would not be bad in the first place. In fact, these deaths would be good. That is hard to believe.

²⁴ For more on the distinction between wide and narrow proportionality, see J McMahan, *Killing in War* (Oxford University Press 2009) 21.

²⁵ For clarification and development, see T Hurka, "Proportionality in the Morality of War" (2005) 33 *Phil & Public Aff* 34; M A Newton and L May, *Proportionality in International Law* (Oxford University Press 2014).

The desert-adjusted theorist might here dig in her heels. Once we fully appreciate that people are getting more than they deserve, she might protest, we should welcome the result that wide proportionality is satisfied. For example, if we focus on the viciousness of the noncombatants in *Eastern Civilians*, the verdict delivered by the desert-adjusted account is intuitive.

This protest, however, does little to help. We need not postulate that the noncombatants are vicious *at all*. Two variables determine where we are on our desert-mountains: overall welfare and desert levels. And remember that the desert-adjusted account, as a whole-life approach, is atemporal. Suppose, then, that the noncombatants in *Eastern Civilians* happened to be completely virtuous but extremely fortunate. Given the sheer magnitude of welfare their lives already contained, they are on their eastern slopes. J is thus not prevented by wide proportionality from bombing these virtuous noncombatants. That is even harder to believe.

The problem made vivid by wide proportionality, as should now be clear, is traceable to *Peaks*. Past the point at which people are getting their due, additional welfare makes things worse. Accordingly, after this point, the bad usually associated with killing vanishes. The deaths of those on their eastern slopes carry no weight. The potential noncombatant deaths in *Eastern Civilians* therefore do not factor positively into the wide-proportionality calculation, as it seems they should.

One response to this problem relies on dropping an assumption we've been making: namely, that the desert-adjusted account treats the fit between desert and welfare as exhausting intrinsic goodness. Call this the *complete axiology approach*.²⁶ Though he does not commit either way, Kagan flags the possibility of taking a different tack. As he explains:

[I]f we accept a pluralistic theory of the good, where other factors beyond desert have intrinsic value as well, then it might well be the case that in terms of one or more of these *other* factors, more good overall is done by having one person be at their peak rather than another.²⁷

In adopting this strategy, one could hold that welfare contributes to intrinsic goodness regardless of where the welfare recipient is on her desert mountain. On this *incomplete axiology approach*, one factor would be one's desert-adjusted welfare and another, independent, factor would be one's welfare. Both of these factors would hence need to come into play to determine total intrinsic goodness. On such an approach, even if killing U's noncombatants would be good from the point of view of desert, it might still, given their loss of future welfare, be bad overall.

An incomplete axiology makes the desert-adjusted account hard to pin down. If other factors can come into play, then it is difficult to tell if a problem we identify isn't remedied by one of these other factors.²⁸ But we can still say this: If the desert-adjusted account is true, it is (at least to some extent) better if a person east of their

²⁶ For example, Feldman (n 13) presents the desert-adjusted account as a complete axiology.

²⁷ Kagan (n 1) 141. McMahan (n 2) ch. 2, for example, takes the incomplete approach—adjusting for desert is only part of his total axiology.

²⁸ Nor do we know how these factors interact.

peak does not receive additional welfare. In other words, there is, even on this incomplete approach, *some reason* to prevent further welfare from flowing to this person. Perhaps this reason is outweighed by other factors. Nonetheless, it is worse, according to the desert-adjusted account, when people continue to move east of their peaks.

We can test this claim. But we need a way of isolating the reason that adjusting for desert provides. To build such an isolation test, it will help to have an example. Consider a narrow-proportionality case

Eastern Army: Country U is an unjust aggressor, invading J's abandoned northern territories. J cares little about these territories. Indeed, suspicious of U, J has rigged these territories with landmines. The landmines can be remotely detonated. Detonating is the only way for J to gain back the north. If J detonates the landmines, thousands of U's troops will be killed. Reliable reports reveal that U has no further plans to press the invasion, and the lives of all U's combatants are, given their desert levels and previous welfare levels, east of their peaks.

To make the isolation test work, we need to make a somewhat contrived stipulation: the loss in future welfare of U's combatants is exactly the same amount of welfare that J would lose in giving up the northern territory. This isn't as gimmicky as it might first appear. Imagine that all of U's combatants will live lives that, due to ups and downs in welfare, are just barely positive. Next, imagine that J's citizens care some, but not much, about these lands. If J gives up the northern lands, J's citizens would lose some welfare. The loss in territory thus brings with it negative welfare, but just barely.

Eastern Army is admittedly stylized, but our isolation test is straightforward. To see how it works, we need to fix some terms. An action is *impermissible* if and only if refraining from the action is required. An action is *merely permissible* if and only if it is an action that is permissible but not required; we have sufficient (but not decisive) reason to act in this way. An action is *required* if and only if there is decisive reason to perform the action; it is the uniquely permissible action.²⁹

To repeat, for an action to be *merely permissible* (permissible but not required), there must be sufficient but not decisive reason to perform the action. This is where the isolation test comes in. Suppose we want to know if something provides us with a reason to act in a certain way. If we know, independently, that some act is merely permissible—for example, killing in self-defense—then we know that any additional reason tips the balance. A reason in favor of killing would make killing required; a reason against killing would make killing impermissible.

Suppose that, given U's unjust invasion, it would be merely permissible for J to detonate the landmines. J's detonating satisfies narrow proportionality; killing U's combatants falls within the scope of their liability. So, on the basis of U's act of aggression, there exists sufficient, but not decisive, reason to detonate. Next, remember that the welfare levels previously experienced by U's army are such that additional future

²⁹ Defining deontic verdicts in this way is popular and compelling. See, for example, D Parfit, *On What Matters*, Vol 1 (Oxford University Press 2011) 32; M Schroeder, "What Makes Reasons Sufficient" (2015) 52 *Am Philosophical Quart* 159; and most importantly for our purposes S Kagan, *The Limits of Morality* (Oxford University Press 1989).

welfare would reduce intrinsic goodness. This badness supplies J with another reason to detonate. Thus, according to the desert-adjusted account, J is *required* to detonate. That is not so.³⁰

The incomplete axiology approach makes desert-adjustments' contribution to goodness hard to discern. Given that the point here is subtle, another pass is warranted. In *Eastern Army*, we stipulated that the lives of U's combatants contained more welfare than deserved. This surplus of welfare puts U's combatants on the eastern side of their peaks. Any extra welfare, accordingly, will move these combatants farther down their eastern slopes, making things worse. Next we imagined a situation where, for reasons unrelated to the goodness or badness of their deaths, it is merely permissible to kill U's combatants. The unjust aggression got us most of the way there. But, since U's combatants had positive future welfare, we needed the contrived stipulation to ensure the balance of reasons. With things balanced, we could run the isolation test. Intuitively, our deontic verdict should stay put at permissibility—detonating is optional. But the desert-adjusted account gives J an additional reason to detonate. If J abstains from detonating and U's combatants receive more welfare by continuing to exist, a worse state of affairs would obtain. This reason tips the balance. J is *required* to detonate. This result, however, strains credulity. The desert-adjusted account fails our isolation test. It tells us there is a reason where there is not one. The account thus needs to be either revised or abandoned.

As we have seen, *Peaks* has troubling implications. But the counterexamples above rely on putting combatants and noncombatants on the eastern side of their peaks. A simple revision is thus available. Abandon *Peaks* in favor of

Plateaus: The line for each level of virtue is plateau-shaped, with a highest point—the kink—followed by a straight line. This kink is what each person, in terms of overall welfare, deserves.³¹

Plateaus tells us that there is a certain amount of total welfare one deserves. It is bad if one gets less than her due. However, additional welfare beyond what one deserves, unlike *Peaks*, does not make things go worse. We can represent the *Plateaus* variant of the desert-adjusted account by revising Figure 12.2.

As is clear from Figure 12.3, *Bell Motion* and *Shift* are still in effect. With the presence of *Plateaus*, however, *Shift* ensures that Becky deserves more welfare than culpable Andy. And *Bell Motion* ensures that, if both are below their kinks and we have equal amounts of welfare to distribute, it is better for it to go to Becky.

The adoption of *Plateaus* mitigates the problems introduced by having more welfare than deserved. In *Eastern Army*, for example, J no longer has a reason to kill. Thus it is

³⁰ Of course, there are different accounts of narrow proportionality: see D Rodin, *War and Self-Defense* (Oxford University Press 2002); McMahan (n 24); and S Uniacke, "Proportionality and Self-Defense" (2011) 30 L & Phil 253. And there are different accounts of the liability to defensive harm: see J J Thomson, "Self-Defense" (1991) 20 Phil & Public Aff 283; M Otsuka, "Killing the Innocent in Self-Defense" (1994) 23 Phil & Public Aff 74; McMahan (n 3); and J Quong, 'Liability to Defensive Harm' (2012) 40 Phil & Public Aff 45. *Eastern Army* will need slight modification depending on one's particular commitments. I trust making such modifications will not prove demanding to the point of distraction.

³¹ For further development of *Plateaus*, see Kagan (n 1) ch. 5.

Fig. 12.3 Desert-Adjusted Account with Plateaus

permissible, but not required, to detonate. The desert-adjusted account with *Plateaus* passes our isolation test.

Still, *Plateaus* faces problems. It tells us that if people are east of their kink, movements east along the X-axis are neutral. Accordingly, though things do not get worse by having an abundance of welfare, they also do not get worse if this surplus welfare is lost. Until, of course, we hit the kink.

Return to *Eastern Civilians*. Suppose that a number of noncombatants have welfare and desert levels far beyond their kink. On this supposition, a deprivation of the welfare from the current position west to the kink is not bad. For now, treat the desert-adjusted account with plateaus as taking the complete axiological approach. Imagine that the deprivation of welfare caused by the deaths of the noncombatants is less than the welfare from their current position to their kink. On this complete approach, given *Plateaus*, wide proportionality would again not rule out J's bombing of the non-combatants. The deaths would not be bad. That is the wrong verdict. But we should not be hasty. The desert-adjusted account, as an incomplete axiological theory, does not deliver this implausible result.

We can now take stock. We've considered *Peaks* and found it delivers implausible verdicts for *in bello* and *ad bellum* proportionality calculations. *Peaks* should be abandoned. *Plateaus* is an improvement on *Peaks*. Yet, combined with the complete axiological approach, the desert-adjusted account still yields implausible results. We should reject the desert-adjusted account as a complete axiology. But the desert-adjusted account with *Plateaus*, as an incomplete axiology, still stands. So our residual question is this: Should we accept this more modest version of the desert-adjusted account when weighing lives? We should not.

4. The Problem with Shift

As an incomplete axiology, the desert-adjusted account with *Plateaus* is dialectically slippery. If everyone gets more than they deserve, this is not bad. And if welfare is another feature in the axiology, then getting more welfare is intrinsically good.

That is, from the point of view of desert, this extra welfare would be neutral. But, since welfare by itself is good, this additional welfare would make things go better overall. Despite this slipperiness, we can test the account by turning to another principle: *Shift*.

Shift, recall, tells us that the more deserving deserve more welfare. As one's virtue goes up, one's kink moves east along the X-axis. As one becomes more vicious, one's kink moves west. Notice that if a person is below her kink, *Shift* tells us that there are two routes to making things better. First, this person can receive more welfare. Second, and more interesting, this person can become more vicious. By becoming more vicious, this person thereby moves her kink west. Less welfare is required for things to be going as well as they could.

Fig. 12.4 Pre-Shift

Fig. 12.5 Post-Shift

As Figures 12.4 and 12.5 illustrate, welfare levels stay static. The dot remains fixed along the X-axis. But, by moving the plateau west, intrinsic goodness goes up. So, if the desert-adjusted account is true, then, assuming we cannot adjust welfare levels, we are given (at least some) reason to shift west. After all, things would be better. Since welfare levels can stay fixed, we can test this claim even if the desert-adjusted account adopts the incomplete axiology approach.

To see *Shift's* troubling implications, consider

Western Siege: Country J is fighting a just war with Country U. Sadly, to bring the war to an end J must lay siege to U's capital, making things go significantly worse for U's citizens. J cannot raise the welfare levels of U's citizens, but can provide vicious-making reeducation pamphlets. J has spent considerable resources ensuring that if U's citizens follow the pamphlet's instructions they will become, given their current low levels of welfare, exactly as vicious as needed to make things as good as they can be.

According to the desert-adjusted account, J has reason to supply the pamphlets, and U's citizens have reason to follow their instructions. Such a shift would promote the overall level of intrinsic goodness. But that is hard to believe. J has no reason to offer the opportunity to make people more vicious, and U's citizens have no reason to take up this offer. Insofar as the desert-adjusted account tells us otherwise, it should be rejected.³²

Of course, many believe that matters beyond your control cannot influence how deserving you are. If we accept this control condition, it might be objected that U cannot cause J's citizens to become more vicious. For, in that case, J's citizens would not, in fact, be less deserving. But a control condition will not blunt the force of the objection raised here. U is not causing J's citizens to become more vicious. U is merely offering an avenue for J's citizens to increase the overall level of intrinsic goodness. The citizens of J can opt to read and implement the pamphlets' instructions or not. It is up to them.

Shift gives rise to a damning problem: Things can go better by people becoming more vicious. Yet, unlike the problem with *Peaks*, no non-arbitrary revision is open to the proponent of the desert-adjusted account. But *Shift* cannot simply be jettisoned. If abandoned, the desert-adjusted account would be robbed of the idea that the more deserving deserve more—a devastating result. Indeed, without *Shift*, the account could not distinguish different virtue levels at all. So long as *Shift* is part of the account, however, the problems made vivid by *Western Siege* persist. The desert-adjusted account should be abandoned.

5. Conclusion

Rejecting the notion that we should adjust for desert when weighing lives is a significant result. The desert-adjusted account provides a powerful and popular rationale for treating just and unjust lives differently. In rejecting this account, we reject the best justification for weighing the lives of the just and the unjust differently. Absent some alternative rationale, assuming equal future welfare, we should weigh the lives of just and unjust combatants equally.

³² Welfare levels in *Western Siege* are fixed. So assume that, in becoming more vicious, citizens do not reduce overall welfare levels by their viciousness. If that is difficult to believe, imagine that Country U offers a mixed package: viciousness reeducation pamphlets, and welfare to offset any loss in welfare caused by J's citizens being more vicious. Though this revision makes the case more complicated, the unintuitive implications of *Shift* should still be obvious.

As seen above, the desert-adjusted account, if true, would have important consequences for just war theory. So too would the straight-welfare account. But I'll leave defending the straight-welfare account and its implications for another occasion. Instead, I want to conclude by taking a pass at something left hanging from Section 3. The lingering worry is this: If the desert-adjusted account is false, why is it so appealing to treat the vicious differently than the virtuous in *Single Death* and *Factory*?

The answer, I suggest, is that we bring about better states of affairs if we act *as if* the desert-adjusted account is true. By having desert-sensitive norms, we incentivize non-culpable behavior and disincentivize culpable behavior. Desert helps us get along given the world we live in and the creatures we are. It directs our behavior, supplying incentives and disincentives, when natural inclination isn't up for the task. In prisoner's dilemmas, for example, desert serves to change the payoff matrix so as to make coordination rational. This is important because human creatures are especially deficient. People are good at looking after their own interests. No system of sanctions is needed to ensure that people attend to themselves. Rather, we need to correct for our tendency to attend too closely to our own interests, either at the expense of others, or at their neglect. It is then unsurprising that we deploy desert, as a useful fiction, to facilitate this end. Thus the trouble with relying on cases like *Single Death* and *Factory* is that we have strong reasons, not provided by the *intrinsic* normative significance of desert, for treating the just differently from the unjust.

If this debunking explanation of our intuitions is correct, although it may appear otherwise, desert does not adjust the intrinsic goodness or badness of a death. Instead, when it comes to desert, we are prone to fetishizing: mistakenly treating what is merely of instrumental significance as if it has intrinsic significance. We should avoid this mistake. Lives, after all, may hang in the balance.