

Oxford Studies in Normative Ethics

Volume 3

EDITED BY
MARK TIMMONS

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 2013

The moral rights of the authors have been asserted

First Edition published in 2013

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

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: 2013938356

ISBN 978-0-19-968590-5 (Hbk.)

ISBN 978-0-19-968591-2 (Pbk.)

Printed and bound in Great Britain 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.

9	They Can't Take That Away From Me: Restricting the Reach of Morality's Demands	203
	SARAH STROUD	
10	What We Know and What We Owe	235
	VANESSA CARBONELL	
11	Objective Double Effect and the Avoidance of Narcissism	260
	HOWARD NYE	
12	Some Intellectual Aspects of the Cardinal Virtues	287
	PAUL BLOOMFIELD	
	<i>Index</i>	315

Objective Double Effect and the Avoidance of Narcissism

HOWARD NYE

The Doctrine of Double Effect [DDE] states roughly that it is harder to justify causing or allowing harm as a means to an end than it is to justify conduct that results in harm as a side effect. The DDE is typically interpreted as maintaining that there are stronger moral reasons against causing or allowing harms with the intention of doing so than there are against causing or allowing harms that we foresee but do not intend. Let us call any such reading of the DDE, according to which our moral reasons and permissions depend upon our intentions, a *subjective reading*.

Over the past several decades, various authors have argued convincingly that our moral reasons cannot depend upon our intentions in the way that subjective readings allege.¹ It is, however, possible to interpret the DDE as maintaining that it is harder to justify causing or allowing harm as a means because of something about the objective explanatory relationship between the effects our conduct has on those it harms and the effects it has on those it benefits. Such *objective readings* have, however, been largely dismissed as absurd.²

In this essay I defend an objective reading of the DDE. I argue that a theory of deontological constraints on harming needs something like the DDE in order to avoid the charge that it reflects a narcissistic obsession with our personal moral purity instead of an appropriate concern for the welfare of others. But, I contend, the central problem with subjective readings of the DDE is that, by making our own intentions more

¹ See especially Thomson (1991, 292–6) and Scanlon (2008). See also Ross (1930, 4–6) and Bennett (1981, 96–8; 1995, 194–6).

² See Frey (1975, 279–83), Bennett (1995, 198–9), and Norcross (1999, 115–17).

important than the welfare of others, they embody exactly this kind of implausible narcissism.

I believe that objective readings of the DDE have been dismissed primarily because they have been construed as claiming that the explanatory relationship between an act's harmful and beneficial effects is a reason against performing it. But I suggest that we should instead interpret the DDE as a denial of the Machiavellian dictum that the ends justify the means. On this reading, the DDE holds that the benefits of our conduct for some individuals do not count (as strongly) in its favour if they come at the expense of others. This, I argue, makes best sense of the original motivations for the DDE and provides a plausible foundation for deontological constraints.

11.1 WHY DOUBLE EFFECT?

The DDE is often invoked as a justification of two kinds of intuitions about particular cases. The first are intuitions about the *permissibility of collateral damage*. Many intuit that it could be permissible to engage in tactical bombing that one knows will destroy enemy military installations and kill civilians as a side effect, yet wrong to target civilians in a terror raid that one knows will result in identical benefits and harms. Many also intuit that it would be permissible to divert a trolley from a main track on which it will hit five people onto a side track where it will hit one, but wrong to stop a trolley from hitting five people by pushing a fat man into its path. Since the terror bomber and fat man pusher seem to cause harm as a means while the strategic bomber and trolley diverter seem to cause harm as a side effect, the DDE appears to provide a natural justification of these intuitions.³

The other kind of intuitions the DDE is invoked to justify concern certain kinds of *wrongful omissions*. It seems permissible to withhold a supply of life-saving drugs from one individual who needs the entire supply in order to give it to five others who each need only 1/5.⁴ But it seems wrong to withhold life-saving drugs from an individual infected

³ See Foot (1967, 23–4), Costa (1986), Shaw (2006), Bennett (1981, 95), and Quinn (1989b, 336).

⁴ Or at least to flip a coin to decide what to do (Taurek 1977). In what follows, Taurekians can replace my talk of the permissibility of benefiting a greater number with talk of the permissibility of flipping a coin.

with an old strain of a disease in order to observe its fatal progression and learn how to cure five people infected with a newly mutated strain. A plausible justification of these intuitions is that failing to give the drugs to the one allows harm to her as a means of saving the five in the second case (where they would not be saved unless her disease progressed), but as a side effect of saving them in the first (where her affliction plays no role in saving them).⁵

I am, however, quite skeptical of our intuitions about permissible collateral damage. If the salvation of five individuals can justify inflicting harm on one as a “side effect,” we would seem justified in driving over one trapped on a road if that was the only way to save five from drowning, but this seems about as abhorrent as pushing the fat man into the trolley’s path. To many it also seems wrong to save five by performing a surgery that will release lethal fumes into a room in which one is trapped, or by destroying a trolley headed towards them with an explosion that you know will envelop a bystander.⁶

Moreover, I do not believe that we should in any event treat our intuitions about what it is wrong or permissible to do in particular cases as a kind of data that our moral theories must fit. Our moral judgments are subject to a host of distortionary factors, such as indoctrination, the asymmetric salience of different considerations, and confusions in understanding and reasoning. The only way in practice to determine whether our convictions reflect these biases is to determine whether they can be supported by general ideas and principles that are directly plausible, or seem true independent of inference.⁷ Our reasons to accept general principles have more to do with whether their plausibility can survive critical clarification and integration with other plausible ideas than whether they match our pre-theoretical intuitions about cases.

⁵ See Foot (1967, 24–5) and Quinn (1989b, 336).

⁶ See Foot (1984, 179, and 1967, 29). Principles have been suggested that would permit “diverting threats” without permitting these acts (see Thomson 1976, 216–7—who has (2008) renounced her proposal—and Kamm (2007, 147). But these principles attribute intrinsic moral relevance to such factors as the identity of the material with which you harm someone (was it the same as that which would have harmed the five?), how close it was to her (was it “in her context?”), and how “directly” it harmed her, which—like your victim’s skin color—seem *obviously* devoid of intrinsic moral relevance. It thus seems to me highly unlikely that our intuitions about permissible collateral damage reflect defensible moral ideas. I suspect that they reflect instead the greater salience of the good we are doing in relation to the harm we are inflicting, and the fact that in real life acts that risk harming as a side-effect are more likely to secure goods and less likely to inflict harms than acts of inflicting harm as a means.

⁷ See Singer (1974, 515–17).

What is most important about the DDE is, I believe, that it is essential to a defensible theory of deontological constraints on harming.⁸ If there are any deontological constraints on harming, I take it that they must explain why it is permissible to do things such as saving five drowning swimmers rather than one, but wrong to push a fat man in front of a trolley to save five. Perhaps the most intuitively obvious account of this is the Doctrine of Doing and Allowing [DDA], according to which there are stronger moral reasons against inflicting harm than there are against failing to prevent harm.⁹ To many of us, this seems directly plausible.

But there is a major threat to the plausibility of the DDA. As Kai Nielsen (1972, 330) suggested, it can seem to be telling us simply to retain our own “moral purity” by avoiding “dirty hands.” Consider:

Alastair and the Fat Man. A trolley is headed toward five people. You see a utilitarian named Alastair sneaking up on an oblivious fat man, ready to push him onto the tracks in order to stop the trolley. You can keep quiet while Alastair pushes the fat man, or call out to warn him of Alastair’s approach.

By itself, the DDA seems to entail that you are forbidden to push the fat man yourself, but allowed, out of concern for the five, to keep quiet and let Alastair do the dirty work. It seems to permit you to benefit the five at the fat man’s expense by staying out of Alastair’s way, just not by taking Alastair’s place. As such, the DDA threatens to embody a narcissistic obsession with your personal purity instead of an appropriate concern with how you treat others. Call this *the dirty hands objection*.

To avoid the dirty hands objection it seems that a theory of deontological constraints on harming must explain why it is wrong, not only to push the fat man, but to let Alastair do it. You might think that this is because Alastair’s pushing would be wrong, and you would be complicit in this wrong if you fail to prevent it. But this seems incorrect. First, suppose that you could save 1,000 people from being killed by a murderer or 1,001 from being killed by a natural disaster. If failing to prevent wrong

⁸ That is, a theory according to which certain ways of causing or allowing harm are in themselves or intrinsically harder to justify than others, in a way that does not simply reflect which acts or policies would bring about the best states of affairs.

⁹ Where the reasons against inflicting harm are *sufficiently stronger* that they decisively outweigh the reasons to save five. When I speak of ‘the DDA’ and ‘the DDE’ I have in mind versions of this kind.

was so much worse than simply failing to prevent harm, you would be required to save the 1,000, but surely you are permitted to save the 1,001.¹⁰ Second, it would not make a moral difference if Alastair were replaced by a humanoid robot, a giant Roomba, or even a moving steel rod. The claim that it is permissible to let these things do the dirty work of pushing the fat man onto the tracks but wrong to do it yourself is equally vulnerable to the dirty hands objection.

It appears that a plausible account of the why it is wrong to let something push the fat man must make reference to something like the fact that you would be passively using his death as a means of saving the five. It must, in other words, invoke the basic idea of the DDE.¹¹

There is, of course, a famous problem with the idea that in saving the five by allowing the fat man to be pushed, you use his death as a means of saving them. For what saves the five is simply the collision of the trolley with his body; the fact that this injures him, and certainly the fact that he suffers the harm of death some moments later, plays no role in their salvation. As several authors have pointed out, harm itself is almost never a means in cases to which the DDE is applied; what is a means is only something intuitively “close to” harm.¹² The *problem of closeness* for proponents of the DDE is to give a plausible and principled clarification of the

¹⁰ See Quinn (1989b, 347) and McMahan (2009, 358). To explain why it is *wrong* to let Alastair push the fat man, reasons against allowing wrongful death would have to be so important as to make it *wrong* to prevent four more non-wrongful deaths instead (and thus certainly wrong to prevent only *one* more non-wrongful death instead).

¹¹ In his attempt to justify intuitions about wrongful omission without the DDE, Scanlon (2008) considers exactly one very special kind of case, where to save five you must allow a victim to die, after which others will harvest his organs. Scanlon simply *assumes* we have an obligation not to take (or, evidently, allow the taking of) a *living* person's organs to save five, and notes that it would be crazy to think “the advantages of our being relieved of this obligation by his dying... justify an exception to the principle requiring us not to kill that person, or to save [his] life when we can easily do so” (33–5). But the question is *why*, if we could save five others by not saving the single person, there *is* any principle “requiring us... to save [his] life when we can easily do so.” There is no principle that requires this in the case of saving five drowning swimmers rather than one, so why would there be such a principle in the case of saving five by not saving the one whose organs will be used to save them? The DDE (unlike Scanlon) offers a principled answer to this question, which applies to other cases of wrongful omission, like allowing someone's disease to progress or allowing her to be pushed in front of a trolley, where Scanlon's remarks about “allowing someone to die to relieve ourselves of an obligation we have while she is alive” are obviously inapplicable.

¹² See Foot (1967, 21–22), Bennett (1981, 98–116; 1995, 201–13), and Quinn (1989b, 336–44).

Doctrine according to which it is difficult to justify using as a means not only harm but the relevant sort of thing “close to” harm.¹³

There are several proposals about how to solve the problem of closeness, a detailed discussion of which is beyond the scope of this essay. For now I wish only to indicate that (1) anyone who wishes to give a plausible response to the dirty hands objection is in the business of solving the problem of closeness, and (2) the problem is not utterly hopeless. To answer the dirty hands objection we must explain why it is wrong to allow something to push the fat man into the path of the trolley to save five. But to explain why it is wrong to do this it seems that we must appeal to the idea that this would involve something like harmfully using the fat man as a means, saving the five at the fat man’s expense, or sacrificing the fat man to save the five.

The best solutions to the problem of closeness identify the DDE’s root idea with the claim that it is particularly difficult to justify treating individuals in one of the foregoing ways, and to seek to make the idea precise. Warren Quinn’s (1989b) “harmful involvement” solution is an excellent example. Quinn argued that we should interpret the DDE as claiming that it is particularly difficult to justify harmfully using someone as a means. His precise suggestion was that we harmfully use someone just in case we involve her in our plans—or use as a means her instantiation of some property—which involvement or instantiation in fact causes her to be harmed. Thus, when you save five swimmers rather than one, the one’s instantiation of properties that harm him (*being in the water, slipping beneath the waves*) are completely immaterial to your salvation of the five. But when you save five from the trolley by allowing the fat man to be pushed in front of them, you use as a means the trolley’s striking him, and this does in fact harm him.

¹³ Some authors (Fitzpatrick 2006, Shaw 2006) seek to solve the problem of closeness by maintaining that in the relevant cases events that constitute harm would be means, so harm itself is a means. It is, however, preposterous to maintain that the event of the fat man’s dying is identical to or constituted by the event of his being impacted with the trolley. The former could take place several minutes or hours after the latter and at a completely different location (if he were rushed to hospital). And it is preposterous to claim that the moral barriers to using as a means a “harm” like mere violent impacting, quite independent of death, are anything like the moral barriers to using death as a means. If you could quickly anesthetize the fat man before pushing him and, after the trolley had hit him, quickly reassemble his body so he awoke a few minutes later without noticing a thing, I venture that everyone should agree that you would be *required* to do so.

An alternative solution can actually be derived from what Quinn (1989a) misleadingly characterized as a version of the DDA. On this view, it is particularly difficult to justify not only actions that produce harm, but deliberate failures to prevent “actions of objects or forces over which we have control” that produce harm.¹⁴ Quinn’s idea seemed to be that it is difficult to justify benefitting some by deliberately producing or allowing events that produce harm to others.¹⁵ Thus, when you save five swimmers rather than one, the events that produce the one’s death (his remaining in the water, his slipping beneath the waves) play no role in saving the five, so you need not intend them. But saving the five on the track by allowing the fat man to be pushed requires events that produce his death (his moving in front of the trolley, his being struck by it), which you must intend as means. While Quinn presents this as a version of the DDA, it does not seem that in allowing the fat man to be pushed you inflict harm on him in any familiar sense. Quinn (1989a, 300) notes that in deliberately allowing an event that produces harm, your agency seems implicated in the harm, but this seems more like a Double Effect idea than a Doing/Allowing idea. Indeed, I think it is plausible that our intuitive distinction between:

- (i) benefitting some at the expense of (or by sacrificing) others, and
- (ii) simply benefitting some instead of others

can be made precise by something like Quinn’s distinction between

- (i’) saving some in virtue of ensuring the existence of events that produce harm to others, and
- (ii’) saving some simply in virtue of failing to prevent events that produce harm to others.¹⁶

¹⁴ Ned Hall (2004) suggests that there are at least two concepts of causation: counterfactual dependence and “production.” He characterizes “production,” in the actual world, as obeying transitivity, locality, and intrinsicness; perhaps it is something like conserved-quantity transfer (Dowe 1995) or trope-persistence (Ehring 1997). I use “production” to refer to what Hall would call “production *and* dependence.”

¹⁵ More precisely, it is difficult to justify “most proximate contributions” or effects of one’s conduct on the whole that have these effects (Quinn 1989a, 301–2).

¹⁶ There are, however, reasons to broaden (i’) to include saving some in virtue of ensuring certain events that result in harm through what Hall (2004) calls “double-prevention.” Whether certain acts that double-prevent harm are morally akin to inflicting harm is debated (see McMahan 1993 and Hanser 1999). A satisfactory specification of (i’) would require a resolution of these difficult issues.

Since a plausible root idea of the DDE is that it is harder to justify (i) than (ii), precisifying this as a moral distinction between (i') and (ii') seems like a promising solution to the problem of closeness.

While these Quinian solutions to the problem of closeness are certainly controversial, they are, I believe, promising enough to provide hope that the problem is tractable.¹⁷

11.2 AGAINST SUBJECTIVE READINGS

I have thus argued that, in order to avoid the dirty hands objection, proponents of deontological constraints must explain why it is wrong to save five by allowing something to push a fat man in front of a trolley, that to explain this we must appeal to something like the ideas of harmfully using or sacrificing the fat man as a means, and that these are best interpreted as “root ideas” of the DDE, precisifications of which will constitute solutions to the DDE’s problem of closeness. I believe, however, that there are at least two sound arguments against the most common understandings of the DDE, which read it subjectively as the claim that there are stronger moral reasons against causing or allowing harmful effects with the intention of doing so than there are against causing or allowing these effects with the foresight but without the intention of doing so.

The first, which I will call *the volitional argument*, runs roughly as follows:

- (P1) What there is stronger or weaker moral reason to do must be something we can voluntarily choose to do.
- (P2) We cannot voluntarily choose to have certain intentions. So we cannot voluntarily choose to perform acts with certain intentions.
- (C1) Therefore, acts performed with certain intentions cannot be what there is stronger or weaker moral reason to do.

¹⁷ For excellent criticism of Quinn’s (1989b) proposal, see Bennett (1995, 218–21). Using Quinn’s (1989a) proposal for the DDA as an understanding the DDE would violate the assumption that there is an intrinsic moral difference between terror raids and tactical raids that are known to have the same consequences (as well as the assumption that it is permissible to divert the trolley—though Quinn mistakenly thought otherwise). But as I indicated in fn. 6, I suspect that no credible principles can support these assumptions.

(C2) Therefore, it cannot be the case that there are stronger moral reasons against causing or allowing harmful effects with the intention of doing so than there are against causing or allowing these effects without this intention (subjective readings of the DDE are false).¹⁸

(P1) seems to follow simply from the practical nature of deontic¹⁹ assessment: what we are asking about in trying to determine what there is moral reason to do is, well, what to *do*: the sort of thing we can choose or will. (P2) also seems obvious, and is illustrated by Kavka's (1983) toxin puzzle, in which an eccentric billionaire will pay you \$1 million if at midnight tonight his completely reliable brain-scanner detects that you have an intention to drink a toxin tomorrow morning which will make you sick for a day. Try as you might, you will *not* be able to form the intention simply in response to the consideration that having it will get you \$1 million. But if intentions were, like movements of our limbs, under our voluntary control, we could form them simply in response to the fact that they will make us rich in the same way that we can extend our arms (to catch \$1 million) in response to the fact that it will make us rich.²⁰

While it is important, I think the volitional argument fails to get to the heart of what is wrong with subjective readings of the DDE. First, the fact that we cannot choose our intentions is probably a contingent fact. We can choose to move our limbs because the neurons responsible for moving them are wired in the right way to the neural correlates of

¹⁸ See Ross (1930, pp. 4–6), Bennett (1981, pp. 96–8; 1995, pp. 194–6), and Scanlon (2008, chapter 2).

¹⁹ *Deontic assessments* of moral reasons for and against doing things, and how they stack up to make acts wrong or permissible, are forward-looking, action-guiding judgments about what to do in a situation. They stand in contrast to *aretaic assessments* that look back upon the quality of an agent's reasoning and motivation in acting, and assign esteem and blame accordingly (see Frankena 1963).

²⁰ I state the volitional argument in terms of the relative strength of moral reasons because that is, I believe, what the DDE is fundamentally a thesis about. But most responses to the volitional argument (McMahan 2009, Wedgwood 2011) have been to versions couched in terms of moral permissibility and impermissibility. By 'permissibility' and 'impermissibility' I mean assessments inextricably linked to moral reasons:

(P3) Something is morally permissible iff it is not decisively opposed by moral reasons, and morally impermissible or wrong iff it is so opposed. So something can be morally permissible or impermissible only if it is something that there can be stronger or weaker moral reason to do.

(P1)–(P3) entail the conclusion directly challenged by critics of the volitional argument:

(C3) Therefore, causing or allowing harmful effects with or without the intention of causing or allowing them cannot be what is itself permissible or impermissible.

voluntary choice. It seems that we can conceive of the neural correlates of voluntary choice being wired to the neural correlates of intention in such a way that we could, on the basis of the good consequences of forming an intention (that it will make us rich in Kavka's puzzle) form it in the same way we can move our arms on the basis of such considerations. But it does not seem that this change in our neural wiring would alter whether a subjective reading of the DDE were true.

Second, the volitional argument leaves open the possibility that something very much like a subjective reading of the DDE is true, namely, a

Modified Subjective Reading of the DDE: There are stronger intrinsic moral reasons against *letting yourself* cause or allow harmful effects with the intention of doing so than there are against *letting yourself* cause or allow harmful effects that you merely foresee.²¹

Even if we cannot make moral decisions about what to intend, if we foresee that we will do something with certain intentions, we can decide to take action to alter those intentions, or decide not to perform the act at all if that is the only way to avoid performing it with problematic intentions. Suppose you are about to save five swimmers rather than your rival, and you suspect that you will do so not only out of concern for the five but in part out of a desire that your rival die. Suppose you then learn that you are being monitored by the Purity Police—a group of demented mind-readers who you know will kill *six* others if you let your rival die with any intention of his dying. Clearly you have very strong reasons in this case to make sure that you do not let your rival die with any intention of his dying; before saving the five you should try to talk yourself into thinking that your rival does not deserve death, or take any mind-altering substances that might remove the intention that he die. If none of this will work, there is a strong case to be made that you should save your rival to minimize the number of deaths.

Modified subjective readings claim that there are powerful *intrinsic* reasons against acting with the intention of harm coming to someone, so we do not need the Purity Police to provide instrumental reasons. On these views it is wrong not to save the fat man because saving him is the only way to avoid allowing him to be pushed with the intention

²¹ Bennett (1995, 195–6) makes exactly this point.

of his being harmfully impacted.²² While this is a coherent position, it seems obviously false. It is plausible that not acting with the intention of someone dying is more important than saving five *if it is the only way to save six*. But how could it be so *intrinsically* important to avoid letting someone die with the intention of her dying that avoiding it *per se* is more important than saving five lives? Suppose you faced the prospect of saving five swimmers or your rival, without any method of purging your propensity to act with some intention of your rival dying if you do not save him, but also without any Purity Police to kill six if you do this. Far from being morally required, saving your rival instead of the five on the grounds that you would otherwise be intending his death looks morally dubious. Even if Taurek is right that you are permitted to save one rather than five, doing so because you would otherwise be intending his death looks like morally bad decision-making. You would be settling a life-or-death question on the basis of a narcissistic obsession with your personal purity rather than an appropriate concern for the welfare of others.

But upon reflection, *unmodified* subjective readings of the DDE seem to face an identical problem. Unmodified readings differ only in that, instead of telling us not to *let ourselves* cause or allow harmful effects with the intention of doing so, they tell us simply not to *cause or allow* harmful effects with this intention.²³ Assuming for a minute that we *can* voluntarily control our intentions, why should it be so morally important that, in doing something that results in harm, we choose to do it without rather than with the intention of a harmful effect occurring? Suppose, for example, that in the foregoing case I saved the five rather than my rival, but I forgot to choose to do it solely out of an intention to save them and ended up choosing to do it in part with the intention of my rival dying. Is this really such a big deal? Is it really wrong in anything like the way failing to prevent something from pushing a fat man into the path of a trolley is wrong?

²² You might also allow this pushing out of sheer indifference, but presumably proponents of modified subjective readings would hold that, so long as you can avoid acting with the problematic intention of someone's being harmed, it is unacceptable not to save someone at trivial cost to yourself.

²³ Wedgwood (2011, 468–9) makes essentially this point. I believe Thomson and Scanlon put their "looking inward" arguments in terms of modified versions because they also accept the volitional argument.

If it is wrong to allow the fat man to be pushed into the path of the trolley to save five (four more than your alternative), it must be wrong to allow him to be pushed in front of it to save two (one more than your alternative). So if what is wrong with allowing him to be pushed is your acting with the intention of a lethal effect on him, it cannot be permissible to choose to act with the intention of a lethal effect on someone to save an additional individual. But consider:

The Impurity Police. You are about to face the situation of saving the five swimmers or saving your rival, and this time you know that your mental states are being monitored, not by the Purity Police, but by the Impurity Police. The Impurity Police credibly promise that they will rescue an additional child from being killed by the Purity Police if and only if you choose to save the five rather than your rival *in part with the intention of your rival dying*.

If it is wrong to choose to act with the intention of a lethal effect on someone in order to save an additional individual, then it must be wrong to choose the option of saving the five with the intention of your rival dying over the option of saving the five without this intention. But surely it is *not* wrong to choose to act with the intention that your rival die in order to save the child. *Who cares* if you have this intention? Surely not your rival, who is going to be allowed to drown either way.

Assuming as we have been that you can choose your intentions, *how could you refuse* to choose to save the five with the intention of your rival dying? How could you explain this to the child's parents? You would have to say: "I'm sorry, but the only way for me to save your child would have been for me to allow someone to die with the intention of his dying. It is true that I allowed him to die anyway. But you see this way I was able to choose to allow him to die without intending it." That would be absolutely monstrous. You would betray the fact that you cared more about the purity of your own intentions than you did about their child's very life.

Of course, you might *try* saying: "It is not that I care about my intentions considered in isolation, but you see if I chose to allow my rival to die with the intention of doing so I would have *disrespected* him, and that is, you know, even worse than just allowing him to come to harm." While this might not be selfish, it is, I think, no less monstrously narcissistic to think that your intentions *per se* have this kind of importance. At this point we should bring in your rival's parents, who should tell

you: “Rubbish. Our son is dead, and you were going to let him die anyway. He never even knew what was in your heart and would hardly have cared if he did know. It had no tangible effect on him whatsoever. Considered in themselves, insults and disrespect are nothing compared to someone’s life. If you could have saved this other child by screaming racial slurs at our son or mocking him as you saved the five it would have been wrong not to do so. Instead, you refused to save this child on the grounds that you would have had this inner state, which our son never even knew about, that was so disrespectful to him that it was more important not to have it than to save this other child? You are seriously sick in the head!”

It seems, then, that both modified and unmodified subjective readings of the DDE are subject to the same central objection: by understanding deontological constraints as moral reasons to be concerned about the intentions with which we act (whether by allowing or simply choosing these intentions), they embody narcissistic obsessions with our personal purity of heart rather than an appropriate concern about what we do to others. I shall call this the *dirty heart objection*. I think the dirty heart objection is particularly damning because, as its name recalls, one of the main attractions of the DDE is its promise to save a theory of deontological constraints from the dirty hands objection that, in forbidding us to push a fat man but allowing us to let other things push him, it embodies a similarly narcissistic obsession with our personal purity.²⁴

11.3 RESISTANCE TO OBJECTIVE READINGS

I have thus argued that a theory of deontological constraints on harming is implausibly narcissistic without the DDE, but that it is also implausibly narcissistic with the DDE given its most common (subjective) reading. One reasonable conclusion would be that there are no deontological constraints on harming. But I do not believe that the DDE’s bid to save deontology from narcissism has yet been given a fair trial. For I believe that there is a better way to understand the Doctrine.

²⁴ The dirty heart objection is, I believe, what gives force to Thomson’s (1991, 291–2) observation that (modified) subjective readings of the DDE implausibly tell us to “look inward” and decide what to do on the basis of the intention with which we would be acting. (Unmodified subjective readings tell us to look inward only to make sure we act with the right intention—but given the dirty heart objection this does not seem much more plausible).

Return to the “root idea” behind the DDE that it is particularly difficult to justify benefitting some at the expense of others. Whether the benefits that some individuals derive from your conduct come at the expense of others is actually an objective, intention-independent matter, about which you could be misled. Suppose you thought that by purchasing Soylent Green rather than cheaper food, you were simply benefitting yourself instead of the children you could have saved by donating the price difference to Oxfam. Conveying such trivial benefits on yourself instead of vitally needed benefits on others is opposed by weighty moral reasons, and I believe that there are conclusive arguments to the effect that it is wrong.²⁵ But suppose you were to learn that (1) Oxfam has been destroyed, and (2) Soylent Green is manufactured by killing children and processing their bodies into the stuff.²⁶ Somehow this seems to constitute a discovery of an even weightier moral case against purchasing it. That such consumption benefits you, not just instead of, but at the expense of children would seem to make your past consumption even more unjustifiable and ceasing consumption even more urgent than you had thought. If Soylent Green were the only food available, it would seem permissible to purchase it if it were not manufactured at the expense of the children. But given that it is so manufactured, there is a strong case to be made that you may not purchase it even if you will thereby starve.

In such a case it would seem crazy to think, as subjective readings of the DDE suggest, that it is unfortunate that you have obtained your new evidence, in the absence of which you could have consumed Soylent Green in peace without intending a harmful effect on the children. If there is something distinctively problematic that you are doing to the children given your evidence, it seems that you were doing it to

²⁵ While Singer’s (1972, 231) talk of “comparable moral significance” and “moral insignificance” can sound unpersuasive or obscure, I believe that his weakened principle is best interpreted as saying something like: “All else held equal, if you can prevent someone from suffering serious harm by incurring only costs that are absolutely trivial in comparison to what she would suffer, it is morally wrong not to do so.” This principle has an *enormous* amount of direct plausibility, and as Unger (1996) has discussed at length, the intuitions it contradicts are extremely dubious.

²⁶ Assume that these children were not going to die at the same time without your help anyway. (Otherwise their being killed might not harm them, or make them worse off than they would have been. See Williams’ “Jim” who can prevent the execution of twenty innocents by killing one of them himself—but see also McMahan’s (2009, 249–52) “altruistic killer”).

them even in your ignorance. Given the unjustifiability of what you were doing, you should be glad to have obtained your evidence, so you can prioritize stopping it at once.

These considerations support an understanding of the DDE according to which it is harder to justify acts which, as a matter of objective fact, benefit some at the expense of others than it is to justify acts that simply benefit some rather than others. Some authors have considered such objective readings of the DDE, but have tended to dismiss them as absurd. There are at least three reasons for this—two of which are not compelling, but one of which is very important.

The first reason some authors dismiss objective readings is that they simply conflate the criteria of objective wrongness and reasons in which these readings are framed with aretaic criteria and criteria of subjective wrongness. Alastair Norcross (1999, 116–17) considers a case such as:

Misleading Evidence. On Friday you have excellent evidence that to your right is one drowning swimmer and to your left are five. So you omit to save the one swimmer and instead pull from the water what you take to be five drowning swimmers. On Saturday there is a party for you at which the Pope gives you a Seal of Approval for gallant action permitted by the DDE. But, just after the Pope awards you the Seal, a hospital official informs you: “You know, it was funny. Those five things you pulled from the water turned out to be convincing inanimate robots. But all was not lost: the organs of the swimmer who drowned turned out to be a unique match for five patients who were dying from organ failure, so we used them to save the five.” Hearing this, the Pope angrily snatches the Seal from you and remarks: “Ah ha! So the benefits your act generated actually depended on the harm it allowed to the one! You should have saved the one instead! I denounce you, and will have no further part in this celebration!”

Norcross’s suggestion is that objective readings of the DDE would entail that the Pope’s obviously inappropriate reaction would be appropriate. But this is false. By way of comparison, consider a case in which you have excellent evidence that you are helping many when you are in fact harming many. Bad criticisms of actual-consequence formulations of consequentialism similarly charge that the theory must be wrong because it entails that in such a case you should have practically reasoned

your way to doing otherwise and that others are justified in blaming you. But as has been pointed out repeatedly, these falsehoods are in no way entailed by consequentialism so understood. Such versions of consequentialism give us a criterion of *objective rightness*, which we are to try to approximate by using our evidence to determine the likelihood that our acts will be supported by the considerations that the theory identifies as reasons.

As we lack omniscience, these subjective assessments of wrongness and reasons in light of our evidence are all we can use in practical reasoning. Whether an agent behaved rightly in the objective but not the subjective sense is completely irrelevant to the quality of her reasoning and aretaic assessments of her blameworthiness or estimability, which are tied to rewards and punishments like denunciations and snatchings of Seals. Objective readings of the DDE, like actual-consequence formulations of consequentialism, are theories of objective moral reasons. It is understood that they will of course be implemented as theories of subjective reasons through reasonable expectations of which courses of action will benefit some at the expense of others. Since in *Misleading Evidence* you did what your evidence told you the objective DDE permitted, it will entail that you blamelessly did right in the subjective sense.

A second reason some authors dismiss objective readings is that they consider versions of the Doctrine that are dubious in ways that are independent of its being read objectively. Frey's (1975, 279–83) early criticism of an objective reading focused on all the implausible features associated with the Catholic tradition: that the DDE prohibits using harmful effects on ourselves as well as others, that its force cannot be attenuated by the wrongdoing or culpability of those who are harmed, and that masturbation is intrinsically immoral. More importantly, I think there can be interactions between the plausibility of objective readings and the ways we assume the problem of closeness should be solved.

It is, I believe, distinctly plausible to say of my *Misleading Evidence* case that because the benefits of not saving the one swimmer came at her expense, the facts of your case did not justify your failure to save her, although you reasonably thought they did. But suppose your evidence told you that the only way to save five innocents was to throw a grenade into a room containing both a sixth innocent and a weapon that

would otherwise be used to kill the five. In fact, the weapon is not in the room and the grenade saves the five by killing the one and consequently demoralizing those who would otherwise have killed them. Here it does not seem to me plausible to say that because the benefit to the five came at the one's expense, the facts of your case did not justify chucking the grenade, though you reasonably thought they did. This, however, is because it seems to me that in saving the five by chucking the grenade, you would have saved the five at the one's expense *even if your evidence had been accurate*.²⁷

A third reason why authors have dismissed objective readings of the DDE is, however, that they have accurately perceived the absurdity of the most simple-minded understanding of what these readings are saying. On this understanding, objective readings are telling us to make sure that our acts do not have beneficial effects that depend upon their harmful effects. That is, they are saying that it is particularly difficult to justify acts that benefit some at the expense of others because the fact that an act benefits some in virtue of causing or allowing harmful effects on others *is itself a weighty reason not to perform the act*. But as Bennett (1995, 199) points out, "There is no evident reason why morality should forbid the [benefit-on-harm dependence] structure" itself.

That is actually an understatement: as Norcross (2008, 76) observes, the view really amounts to the ridiculous claim that it is worse to cause or allow a harmful effect if it does any good. Consider:

The Other Three. You can save two drowning swimmers to your left or one drowning swimmer to your right. You also know that there are three totally different people in hospital dying from organ failure, for whom the organs of the one swimmer are a unique match. So if (but only if) you save the two rather than the one, the one will drown and his organs will be used to save the other three as well.

²⁷ Importantly, Norcross (1999, 116–17) and McMahan (2009, 368–9) consider cases that are more like this than Misleading Evidence. Of course, counting both kinds of grenade chucking as benefiting some at the expense of others makes trouble for the distinction some try to draw between tactical and terror bombing. But as I mentioned in fn. 6, I think harmful tactical bombing *is* difficult to justify, and is typically more easily justified by our evidence than terror bombing only because it is much less likely to inflict harm and much more likely to produce benefits.

On the simple-minded understanding, objective readings of the DDE hold that it is wrong to allow something to push a fat man into the path of a trolley in order to save five because

- (i) the fact that your act will save some in virtue of allowing a lethal effect on others is a reason against performing it, that
- (ii) decisively outweighs the fact that your act will save five.

But this means that although it is permissible to allow one swimmer to drown rather than two where this is all that happens, it is wrong, given the presence of the other three, to allow the one swimmer to drown rather than the two because your act would (i) save some in virtue of allowing a lethal effect on others, which (ii) counts decisively against your act despite the fact that it would save five. It would, in this context, be permissible to let the one drown so long as it does not do any good. This, of course, is preposterous.

11.4 HOW “THE ENDS DO NOT JUSTIFY THE MEANS”

To avoid the dirty hands objection, a theory of deontological constraints must, I have been saying, explain not only why it is wrong to push a fat man in front of a trolley to save five but wrong to let something push him as well. Subjective readings of the DDE say it is wrong to let the fat man be pushed because you would be letting him die with the intention of his dying. But this faces the dirty heart objection that it amounts to a narcissistic obsession with the purity of your intentions. Simple-minded objective readings of the DDE say it is wrong to let the fat man be pushed because your conduct would benefit some individuals in virtue of allowing harm to others. But this seems to entail, preposterously, that it can be permissible to allow harmful effects as long as they do no good.

So why *is* it wrong to let the fat man be pushed into the path of the trolley in order to save five, assuming that it is? It would seem that *the reason it is wrong*, the feature that *makes it wrong*, and your *reason not to do it* is actually the same reason you should not let the fat man be pushed into the trolley’s path when the five are not present: if you allow him to be pushed he will die, and you can easily prevent this. That, I believe, is the most natural and unforced explanation. It appeals to nothing more than his welfare, and the fact that you could easily promote it.

But how could the fact that the fat man will die if you don't save him make it wrong to fail to save him, when the fact that one swimmer will die if you don't save him *doesn't* make it wrong to fail to save *him*? In both cases our reason to save the one seems opposed by the same reason not to save him: that by doing so we can save five others. So how can that reason be a sufficient justification for failing to save the one swimmer but not for failing to save the fat man? The assumption common to both subjective and simple-minded objective readings of the DDE is that this explanation must cite some *additional reason*, beyond just the effect on the fat man's welfare, *against* pushing him into the trolley's path. But this, I fear, is exactly where theories of deontological constraints go wrong.

It is often suggested that in addition to failing to save the fat man, you disrespect him, or somehow offend against his autonomy in a way you do not disrespect or impose upon the swimmer.²⁸ But because your effect on how things are for the fat man is identical to your effect on how things are for the swimmer, it is very difficult to believe that there is any such difference in treatment that could be more important than your reasons to save the five. As we saw with the Impurity Police, it is preposterous that the fat man or his representatives should care significantly about any secret "disrespect" allegedly embodied in your intentions towards him in acting. In the same way, it seems absurd to think that you have offended against the autonomy of the fat man in any way in which you have not offended against the autonomy of the swimmer. In both cases you allow effects that interfere in identical ways with their ability to live their own lives as they see fit.

Of course, if it is *for independent reasons* wrong to allow the fat man to die but not wrong to allow the swimmer to die, then the former or his representatives might justifiably complain of the deprivation of a good where the latter might not, for you owed the good in the one case but not the other. They might even put their complaint in terms of your failing to "respect" the fat man by giving him what you owed him. But, obviously, it is then the antecedently greater difficulty of justifying the failure to save the fat man that explains the disrespect, not the disrespect that explains the greater difficulty of justifying your failure to save him.

²⁸ See Quinn (1989a, 1989b).

How, then, could it be wrong to fail to save the fat man but not the drowning swimmer if there are no moral reasons that count against doing the former that do not equally count against doing the latter? Consider the Machiavellian dictum that ‘The ends justify the means.’ In context, the idea seems to be that if the only way to promote a beneficial end is to use harmful means, the benefit counts as a perfectly good reason to use them. It is, however, plausible to understand deontological constraints on harming as, fundamentally, a rejection of this idea. It is not that we have some special kind of reason *against* using harmful means. It is rather that when an act will benefit some only by having harmful effects on others, the benefits simply *do not count* in the same way as reasons *to* perform it.

Jonathan Dancy has emphasized the distinction between considerations that favour and oppose acts on the one hand and considerations that strengthen or weaken the force of other reasons on the other. For instance, while the fact that you promised to go to the store is a reason to go to the store, the fact that the promise was given under duress weakens this reason without itself counting against going. If you had no other reasons for or against going to the store, the fact that the promise to go was given under duress would move your situation in the direction of both options being permissible; it would not by itself tend to make going to the store something you positively should not do.²⁹

This is, I believe, exactly how we implicitly think about the fact that our conduct will have beneficial effects on some in virtue of its harmful effects on others: it weakens the status of the benefits as reasons to engage in the conduct without counting positively against the conduct. Consider the following pair of cases suggested by McMahan (1994):

Accident Victim 1. An accident victim will die if you do not help him, but your risk of contracting a fatal disease is so great that it is supererogatory to help.

Accident Victim 2. The same as before, except you know that if you fail to help and the victim dies, his organs will be used to save five people in hospital.

²⁹ Dancy (2004, especially 38–52). See also Kagan’s (1988) distinction between features that “additively” make an independent positive or negative contribution to an act’s deontic status and features that “multiplicatively” affect the contributions of other features.

The fact that in the second case the five will benefit at the expense of the accident victim is not a new reason *to help* him. The presence of the five in the second case cannot make it obligatory to help the accident victim where it was supererogatory to do so before. But neither, deontologists should say, does the presence of five in the second case (do much to) add to the case *against* saving the accident victim. What justifies (that is, permits) your not saving the accident victim in the second case is, as in the first, simply the risk to yourself, not the benefit to the five in the hospital. The fact that not saving the accident victim will benefit the five simply does not count (very strongly) in favour of not saving him, because these benefits would come at the expense of the victim.

Along these lines, I suggest that we read the DDE objectively, but as a claim about the *weakening* or undermining of reasons *to* cause or allow harm rather than some new set of reasons *against* doing so.

The Preferred Objective Reading of the DDE: All else held equal,³⁰ the fact that an act or omission will result in benefits for some individuals at the expense of other individuals weakens the extent to which those benefits count in favour of the act or omission.³¹

The preferred objective reading is not saying that there is anything particularly objectionable about the fact that an act has benefits in virtue of having harmful effects; it is not saying we should try to make sure that our acts do not have this property. According to this reading, there is absolutely nothing wrong with saving two swimmers rather than one

³⁰ The all-else-held-equal clause is required because there are plausible factors that attenuate or undermine the applicability of the DDE (these will be weakeners of weakeners). Such factors include consent to be harmed, culpability, and a duty to bear the harm. Candidate attenuators must not be *ad hoc*, but it is, I believe, directly plausible that these considerations undermine the DDE's applicability.

³¹ I speak of the *weakening* of the strength of the reasons constituted by the benefits rather than the *total disabling* of their status reasons in order to allow a non-absolutist formulation of the DDE. Since the DDE is intended to apply to sub-lethal upshots, an absolutist formulation would be intolerable (it would entail that it is wrong to simply push someone down to save someone else's life). Of course, the weakening must be substantial if it is to explain why it is wrong to push or allow the pushing of the fat man, and non-absolutists need a plausible, non-arbitrary account of its degree. For serious harms, I think we should start with the vague idea that the weakening is "massive" or, given the weakening, only "an absolutely ridiculously crazily greater" amount of good could justify the benefit.

when three others will benefit from the one's death by receiving his organs. Here the omission to save the one is fully justified by the fact that it is the only way to save the two. As in a case when the other three are not present, your act simply benefits the two swimmers rather than the one; it does not benefit the two at the one's expense. All the preferred objective reading insists is that because they come at the one swimmer's expense, the benefits to the five in hospital do not count very strongly in favour of allowing the swimmer to drown.

According to the preferred objective reading, allowing the fat man to be pushed in front of the trolley to save the five is wrong, not because its beneficial effects depend upon its harmful effects, but—just like allowing him to be pushed when it does no one any good—because it allows him to die when you could easily prevent it. This is the *reason not to do it* and the fact that *makes it wrong*. The fact that the five would benefit at the fat man's expense merely explains why the benefits to them do not count very strongly in favour of allowing his pushing and why they fail to make this omission permissible. This is why allowing the fat man to be pushed differs from saving five swimmers rather than one. Since saving the five swimmers simply benefits them instead of the one and does not benefit them at the one's expense, there is nothing to prevent the benefits to the five from counting fully in favour of not saving the one and rendering that option permissible.

The assumption that relationships of dependence between beneficial and harmful effects would have to be reasons against acting or wrong-makers is, I believe, the primary reason why people have thought it absurd that they could matter morally. By way of analogy, suppose you made a promise to go to the store that would, under ordinary circumstances, oblige you to go there rather than stay where you are and provide costless help to someone who needs it. But suppose that the promise was made under duress, and that absent these reasons to go it is wrong not to stay and provide help. It is, in particular, wrong to go to the store rather than stay. If someone asked you why it was wrong to go to the store rather than stay, you would not say, "In going to the store, I would be doing what I promised to do under duress." That would be crazy! How could there be anything morally objectionable about doing what you promised under duress to do? Surely the view is absurd on its face! This is not, however, because the fact that a promise was given under duress is morally irrelevant. It is simply because, in answer to the

question of why an act is objectionable, you have cited a consideration that explains why something that could have made it unobjectionable did not, rather than what made it objectionable in the first place.³²

I believe that the preferred objective reading of the DDE gives us a plausible way to rescue a theory of deontological constraints from the charge of narcissistic obsession with the cleanliness of our hands, without running into the charge of narcissistic obsession with the purity of our hearts. Whether you push the fat man yourself or allow him to be pushed, the benefits to the five come at his expense, so according to the preferred objective reading, they fail to count very strongly as reasons to do *or* allow the pushing. Consequently, pushing the fat man yourself or letting something else do the dirty work are wrong for the same reason: they result in harm to the fat man (which you could avoid at trivial cost to yourself). This is a powerfully important fact about the effect of your conduct on the welfare of another individual rather than a dubiously relevant fact about the beauty or ugliness of your internal states.

The preferred objective reading's explanation of why this consideration makes it wrong to let the fat man be pushed, and how this case differs from that of the six swimmers, also looks appropriately focused on how your conduct affects others. On this view, the benefits your conduct would generate for some lose their force as reasons to engage in that conduct, not because of anything about your internal states, but because these benefits would be generated at the expense of others.

³² The fact that the relationship between our conduct's harmful and beneficial effects works as a weakener of the status of its benefits as reasons rather than a reason not to act should help clarify (if it really needs clarification) how this relationship matters morally without a similar relationship between natural events' harmful and beneficial effects mattering axiologically. Clearly, we should not care whether an avalanche kills one rather than five because it simply lands on the one rather than the five or because the one shields the five. Neither outcome is better nor worse than the other (Tadros 2011, 219–20). Of course, this kind of axiological evaluation seems irrelevant to proposals about deontological constraints (as opposed to strange consequentialist views according to which we should save the fat man because it is intrinsically bad for the world to contain instances of some benefiting from the misfortunes of others). But the preferred objective reading helps clarify why this is so: it is a basic, agent-relative deontic fact that the beneficial effects of your conduct do not count (as strongly) as reasons to engage in it if they come at the expense of others. Apart from the harms to those others, no reason against engaging in that conduct—like the alleged intrinsic badness of some benefiting from the misery of others—is needed to explain why it is wrong.

It is important to be clear about how the preferred objective reading makes the reason-giving force of benefits to some dependent upon facts about others. For each individual, we seem to have standing reasons to treat her in general kinds of ways, the strength of which are affected only by facts about her (such as the extent to which the treatment will benefit her) and facts about us (such as whether we are specially related to her as a family member or friend). The preferred objective reading does not maintain that the strength of these general reasons to benefit someone depends on facts about others; what it does is place constraints on the extent to which these general reasons to benefit her can support specific courses of beneficial action or omission. The fact that an act or omission will benefit an individual at the expense of others does not affect the strength of our reasons to pursue the general end of helping her; it merely makes it difficult for this end to justify its pursuit by means of that act or omission. This, I believe, is a directly plausible constraint on our reasons to pursue morally important goals in particular ways, which does not make the moral importance of the general goal of helping an individual implausibly dependent upon facts about others.³³

Because of its appropriate focus on how our conduct affects others, the preferred objective reading enables proponents of agent-centered deontological constraints to give an adequately non-narcissistic justification of their refusal in certain situations to bring about the most good. Suppose a deontologist saves the fat man from falling onto the track, and the parents of the five object: "How could you do this to our children, who otherwise would have been fine? Were their five lives less valuable than the fat man's one?" "Certainly not," the deontologist should reply, "it is just that, under the circumstances, my reasons against allowing the fat man to be pushed outweighed my reasons to bring about the most valuable outcome by means of allowing him to be pushed." The parents inquire: "And what reasons were those?"

At this point, most subjective readings of the DDE would tell our deontologist to make the unacceptably narcissistic reply: "My reasons not to allow the fat man to die with the intention of doing so." But the preferred objective reading allows her to say simply: "My reasons to save the fat man, at trivial cost to myself." "But how," the parents ask, "could those outweigh your reasons to save our *five* children, simply by minding

³³ I am grateful to a referee at Oxford University Press for raising this issue.

your own business?" "Because," our deontologist can explain, "under the circumstances the benefits to your children were not very good reasons against intervening." The parents demand to know "WHY NOT?!" Our deontologist could not look them in the eye and say: "Because to save your children I would have had to allow harm with the intention of doing so." But she need feel no embarrassment in saying what the preferred objective reading entails: "Because the benefits to your children would have come at the fat man's expense; the only reason non-intervention would have saved your children is that it would have ensured the fat man's smashing."

II.5 CONCLUSION

It is quite plausible in the abstract that it is harder to justify benefitting some individuals at the expense of others than it is to justify simply benefitting some individuals instead of others. According to the preferred objective reading of the DDE, this is true because, if doing or allowing something will generate a benefit for some individuals at the expense of others, the benefit loses (much of) its status as a good reason to do or allow that thing. On reflection, I think that there is a *great deal* of direct plausibility to this idea. In fact, it seems to me no less clearly true or self-evident than the Principle of Beneficence that underlies impartial consequentialism, according to which there is intrinsic moral reason to promote the welfare of others. Unless said of situations in which harm to some is the only means of preventing *radically* greater harms to others, the Machiavellian dictum that beneficial ends are perfectly good reasons to use harmful means looks implausible on its face.

While the apparent plausibility of a general ethical principle does not guarantee its truth, we should accept the principle if its plausibility survives critical scrutiny and harmonious integration with other plausible ideas. The best arguments against deontological constraints acknowledge the initial plausibility of something like the DDA or DDE, and attempt to show that critical scrutiny undermines this plausibility. In this essay I have considered what I take to be one of the most powerful such arguments—that, on reflection, deontological constraints seem to embody an implausibly narcissistic obsession with the purity of our hands and hearts. I have argued that this argument fails to undermine an objective reading of the DDE according to which the benefits of our

conduct do not count as strongly in its favour when they come at someone's expense. If I am right about this, and the plausibility of this view survives other forms of scrutiny,³⁴ I believe we should accept it.

Ethical justifications must give out somewhere, and ethical theories need to take certain principles as fundamentally axiomatic or constitutive of the deepest theoretical justifications there are. Like most impartial consequentialists, I think that the Principle of Beneficence is an axiom of this kind. Scrutiny will reveal, I believe, that nothing could be more clearly or basically true than the idea that there are moral reasons to promote the welfare of others. But the idea that the benefits of our conduct for some do not count as strongly in its favour when they come at the expense of others looks to me to be equally axiomatic. Like Beneficence, it does not seem to need any further justification. I think that deontological theories will be on far firmer ground if they acknowledge this idea, rather than anything about respect, rights, or autonomy, as the fundamental axiom underlying constraints on harming.³⁵

REFERENCES

- Bennett, Jonathan (1981). Morality and Consequences. *The Tanner Lectures on Human Values*. Vol. 2. Ed. Sterling McMurrin. Salt Lake City, UT: University of Utah Press, 46–116.
- (1995). *The Act Itself*. Oxford: Clarendon Press.
- Costa, Michael (1986). The Trolley Problem Revisited. *Southern Journal of Philosophy* 24: 437–49.
- Dancy, Jonathan (2004). *Ethics Without Principles*. Oxford: Clarendon Press.
- Dowe, Philip (1995). Causality and Conserved Quantities: A Reply to Salmon. *Philosophy of Science* 62: 321–33.
- Ehring, Douglas (1997). *Causation and Persistence*. Oxford: Oxford University Press.
- Fitzpatrick, William (2006). The Intend/Foresee Distinction and the Problem of "Closeness." *Philosophical Studies* 128: 585–617.

³⁴ Such as whether it can be combined with a convincing detailed solution to the problem of closeness.

³⁵ This essay has benefited from conversations with and feedback from too many people to remember or name. I am particularly indebted to John Ku for many valuable discussions. Special thanks are also due to Tom Beauchamp, Stephen Darwall, Allan Gibbard, Alastair Norcross, Drew Schroeder, Steve Sverdluk, Ralph Wedgwood, and an anonymous referee at Oxford University Press. I am also grateful for helpful feedback from audiences at the Third Annual Rocky Mountain Ethics Congress, the 2009 Eastern Division Meeting of the American Philosophical Association, and the Third Annual Arizona Workshop in Normative Ethics, as well as the students in my fall 2011 seminar on Deontological Ethics at the University of Alberta.

- Foot, Philippa (1967). The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review* 5: 5–15. Reprinted in P. Foot, *Virtues and Vices*. Oxford: Oxford University Press, 19–31.
- (1984). Killing and Letting Die. In J. L. Garfield and P. Hennessey (eds.). *Abortion: Moral and Legal Perspectives*. Amherst, MA: University of Massachusetts Press, 177–85.
- Frankena, William (1963). *Ethics*. Englewood Cliffs, NJ: Prentice Hall.
- Frey, R. G. (1975). Some Aspects to the Doctrine of Double Effect. *Canadian Journal of Philosophy* 5: 259–83.
- Hall, Ned (2004). Two Concepts of Causation. In J. Collins, N. Hall, and L. Paul, Causation and Counterfactuals. Cambridge, MA: MIT Press, 225–76.
- Hanser, Matthew (1999). Killing, Letting Die, and Preventing People from Being Saved. *Utilitas* 11: 277–95.
- Kagan, Shelly (1988). The Additive Fallacy. *Ethics* 99: 5–31.
- Kamm, Frances (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. New York: Oxford University Press.
- Kavka, Gregory (1983). The Toxin Puzzle. *Analysis* 43: 33–6.
- McMahan, Jeff (1993). Killing, Letting Die, and Withdrawing Aid. *Ethics* 103: 250–79.
- (1994). Revising the Doctrine of Double Effect. *Journal of Applied Philosophy* 11: 201–12.
- (2009). Intention, Permissibility, Terrorism, and War. *Philosophical Perspectives* 23: 345–72.
- Nielsen, Kai (1972). Against Moral Conservatism. *Ethics* 82: 219–31.
- Norcross, Alastair (1999). Intending and Foreseeing Death: Potholes on the Road to Hell. *Southwest Philosophy Review* 15: 115–23.
- (2008). Off her Trolley? Francis Kamm and the Metaphysics of Morality. *Utilitas* 20: 65–80.
- Quinn, Warren (1989a). Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing. *Philosophical Review* 98: 287–312.
- (1989b). Actions, Intentions, and Consequences: The Doctrine of Double Effect. *Philosophy and Public Affairs* 18: 334–51.
- Ross, W. D. (1930). *The Right and the Good*. Oxford: Clarendon Press.
- Scanlon, T. M. (2008). *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
- Shaw, Joseph (2006). Intentions and Trolleys. *Philosophical Quarterly* 56: 63–83.
- Singer, Peter (1972). Famine, Affluence, and Morality. *Philosophy and Public Affairs* 1: 229–43
- (1974). Sidgwick and Reflective Equilibrium. *Monist* 58: 490–517.
- Tadros, Victor (2011). Wronging and Motivation. In R. A. Duff and S. Green (eds.), *Philosophical Foundations of Criminal Law*. Oxford: Oxford University Press, 207–27.
- Taurek, John (1977). Should the Numbers Count? *Philosophy and Public Affairs* 6: 293–316.
- Thomson, Judith (1976). Killing, Letting Die, and the Trolley Problem. *The Monist* 59: 204–17.
- (1991). Self Defense. *Philosophy and Public Affairs* 20: 283–310.
- (2008). Turning the Trolley. *Philosophy and Public Affairs* 36: 359–74.
- Unger, Peter (1996). *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.
- Wedgwood, Ralph (2011). Scanlon on Double Effect. *Philosophy and Phenomenological Research* 83: 464–72.