



BRILL



People in Suitcases

Kacper Kowalczyk

Department of Philosophy, University College London,
London, United Kingdom

kacper.kowalczyk@ucl.ac.uk

Abstract

Ex-ante deontology is an attempt to combine deontological constraints on doing or intending harm with the idea that one should act in everyone's interest if possible. I argue that ex-ante deontology has serious problems in cases where multiple decisions are to be made over time. I then argue that these problems force us to choose between commonsense deontological morality and a more consequentialist morality. I suggest that we should choose the latter.

Keywords

deontology – consequentialism – the veil of ignorance – decision theory

1 Introduction

In some cases, acting in everyone's best interest involves doing or intending harm in a way that would otherwise be considered impermissible. Take, for example,

Opaque Footbridge. A, B, and C are trapped in three suitcases. An out-of-control lethal trolley is headed for two suitcases on the track. The other suitcase is on a footbridge above. The suitcases have just been reshuffled, so no one knows who is in which suitcase. You can push a lever to topple

the suitcase now on the footbridge onto the track, stopping the trolley before it hits the other two suitcases.¹

We can represent this case as in Figure 1.




FIGURE 1 Opaque Footbridge

In this decision tree, squares represent your choices, while circles represent moments when your uncertainty is resolved, with different possible resolutions having the probabilities shown. A, B, and C can either die (denoted by a gray or a white skull, depending on whether they are killed as a means or not) or live (denoted by a happy face).²

What should you do?

On the one hand, pushing is in everyone’s interest. The mode of death seems irrelevant to the victim: to be killed as a means seems just as bad for a person as to be killed as a side-effect.³ If you push, everyone has a 1/3 chance of death and a 2/3 chance of survival. If you do not push, everyone has a 2/3 chance of

1 This is a scaled-down version of Hare’s (2016) case. See also Hare (2013: 89–96). Similar cases have been discussed by Thomson (1990: 176–202) and by many others typically under the heading of “survival lotteries.” See Harris (1975) and Singer (1977), but also Kamm (1996: 143–171, 290–310). A real-life “survival lottery,” proposed for Allied bombers during World War Two, is recounted by Glover (1977: 212–213).

2 See McClennen (1990: 99–111) for more on decision trees.

3 Nagel (1986: 178) claims that “it is no worse for the victim to be killed or injured deliberately than accidentally.” See also Scheffler (1994: 109) and Parfit (2011: 365). Scheffler and Nagel credit the point to Scanlon; see Scanlon (1998: 163). Thomson (1990: 184) disagrees, arguing that “dying in consequence of being suddenly grabbed off the street, or out of a doctor’s examining room, and having one’s body organs removed for use to save others” is worse than “dying of organ failure after a long illness.” Still, even if it is in some way bad for a person to be killed as a means, it does not seem so bad that the person cannot be as badly off on balance as if he or she died an accidental death.

death and a 1/3 chance of survival. It is in any person’s interest to face the former prospect rather than the latter.⁴

On the other hand, pushing involves doing or intending harm, which – many people strongly believe – is impermissible except in special circumstances.⁵ These people believe in *deontic constraints* and often appeal to cases like

Transparent Footbridge. A, B, and C are trapped in three suitcases. The situation is the same as in Opaque Footbridge, except that everyone knows who is in which suitcase.⁶

If it is A who is on the footbridge, we can represent this case as in Figure 2.




FIGURE 2 Transparent Footbridge

How can this apparent tension be resolved? I see three main natural theoretical options, as in Table 1.⁷

TABLE 1 Possible views

	Opaque Footbridge		Transparent Footbridge	
	<i>push</i>	<i>-push</i>	<i>push</i>	<i>-push</i>
Permissivism	✓	✗	✓	✗
Ex-ante deontology	✓	✗	✗	✓
Ex-post deontology	✗	✓	✗	✓

4 This follows from a version of Resnik’s (1986: 91–92) better-chance principle, similar to the sure-thing principle but much less controversial.

5 Examples include Nozick (1974) and Kamm (2007). Dissenters include Scheffler (1994) and Kagan (1989).

6 This is a version of Thomson’s (1985) case.

7 Compare Adler and Sanchirico (2006: 289) and Fried (2012: 525–529).

Let us start with the two extremes. The first is what I call *permissivism*. It dispenses with deontic constraints. It is broader than *consequentialism* as it is compatible with the existence of *moral options*: permissions to do less than best.⁸ It implies that you should push in both cases. The other extreme I call *ex-post deontology*. It says that there is a weighty prohibition on doing or intending harm even if that is in everyone's interest. It implies that you should not push in either case.⁹

There are also moderate views.¹⁰ I think that the most attractive one is what I call *ex-ante deontology*. It says that one should choose an option whose prospect is best for everyone (if such an option exists) even if that involves doing or intending harm in a way that would otherwise be impermissible. It implies that, in a Footbridge case, you should push if the case is opaque but should not push if the case is transparent.

Ex-ante deontology can be seen as an attempt to combine deontology with the principle of *Ex-Ante Pareto*, according to which one should choose an option whose prospect is best for everyone (if such an option exists).¹¹ It can also be given a rationale that should appeal to deontologists. To see this,

8 See Scheffler (1994) and Kagan (1989: 183–203) for a view with moral options but without deontic constraints.

9 See Kamm (1996: 290–310) and Alexander (2014).

10 See Thomson (1990: 176–202), Hare (2013, 2016), Frick (2015), and Setiya (2020). I will discuss these authors in more detail below.

11 Ex-Ante Pareto is of course stronger than what is needed to argue that you should push in Opaque Footbridge. It is, at any rate, a plausible principle which also provides a useful reference point. It is plausible because it reflects an attractive individualistic ideal: an option proscribed by Ex-Ante Pareto would seem unjustifiable to the people affected. I lack the space to defend Ex-Ante Pareto against all challenges. But let me mention, first, that many alleged counterexamples to Ex-Ante Pareto will prove spurious once it is combined with an ecumenical account of welfare. (A narrower account of welfare may be suitable for other purposes.) It might, for example, seem impermissible to do a risky but life-saving blood transfusion on a committed Jehovah's Witness even if that affects no one else. (Compare Setiya (2020: 65).) But insofar as I share that intuition, it is because I intuit that the loss of autonomy is a harm. If the loss of autonomy is considerable, then, presumably, the transfusion is no better for the Witness, so this is no counterexample to Ex-Ante Pareto. If the loss of autonomy is slight, then Ex-Ante Pareto's verdict seems correct. (Compare Thomson (1990: 187–191).) Another challenge is presented by Mahtani's (2017) puzzle, which she takes to show that a claim like "everyone's prospect is better if you push" is incoherent as it stands, and, hence, that Ex-Ante Pareto is likewise incoherent. A similar puzzle is discussed by Harman (2015), Hare (2016), and Setiya (2020). I am relatively untroubled by Mahtani's puzzle, however. First, Mahtani's puzzle is similar to a familiar puzzle in quantified epistemic logic (see Gerbrandy (2000)) and, so, we should expect, it is susceptible to a similar solution. Second, Mahtani's puzzle does not threaten some versions of Ex-Ante Pareto – notably, an "objective" version defined in terms of non-epistemic probability – that could be used to argue that you should push in a version of Opaque Footbridge.

recall that Nozick, a prominent deontologist, claims that deontic constraints “reflect the fact of our separate existences,” as the individual harmed “does not get some overbalancing good from his sacrifice.”¹² So, it looks like, for Nozick, deontic constraints are meant to protect people from an uncompensated sacrifice of interests. But, in some cases, doing or intending harm is in everyone’s interest, at least relative to what is known at the time of action. Nozick’s comments suggest that, in these cases, deontic constraints are lifted.

Nonetheless, in this paper, I will argue against ex-ante deontology and other compromise views like it. In sections 2 to 4, I will show that ex-ante deontology has serious problems in cases where multiple decisions are to be made over time. Still, ruling out ex-ante deontology and other compromise views leaves us with a choice between permissivism and ex-post deontology. In section 5, I give a new argument against ex-post deontology’s verdict that you should not push in Opaque Footbridge. In section 6, I also describe a new problem for a deontologist who would like to act in everyone’s interest at least in the limited range of cases where, unlike in Opaque Footbridge, the true state of the world is known in advance. I conclude that it is difficult to combine deontology with *any* Pareto principle, whether ex-ante or ex-post. This means, I suggest, that deontology does not take individual people seriously. That is a serious problem.

2 Problems for Ex-Ante Deontology

Ex-ante deontology runs into problems in cases that are like Opaque Footbridge, except that, after you decide to push the lever, you find out who is in which suitcase and have a chance to change your mind before it is too late. As an example, consider

Case One. A, B, and C are trapped in three suitcases. The situation is the same as in Opaque Footbridge, except for the following details. It is now 12:00. You still do not know who is in which suitcase but, luckily, everyone will peek out of their suitcases by 13:00. The trolley is slow and will only roll under the footbridge at 13:15. The lever is also rusty. You can use it to topple the footbridge suitcase but only if you start pushing right away and keep pushing until the trolley arrives. If you stop pushing at 13:00, then whoever is on the footbridge will see you and will become mildly traumatized, realizing that he or she could easily have died a terrible death as a train stop. Alternatively, you can walk away right now and let the trolley run its course.

¹² Nozick (1974: 33).

We can represent this case as in Figure 3 (with a sad face denoting mild psychological trauma).




FIGURE 3 Case One

What should you do? The bold lines above show what I take to be ex-ante deontology’s answer. At 12:00 it is in everyone’s interest that you push. So ex-ante deontology implies that you should push at 12:00. In terms of the decision tree, this means going up at choice node 1.¹³

But at 13:00 it is no longer in everyone’s interest that you push. If it is A who is on the footbridge, for example, then it is not in A’s interest to push, even if it is still in the interest of B and C. So deontic constraints apply and ex-ante deontology implies that you should not push the lever at 13:00. In terms of the decision tree, this means going down at choice nodes 2, 3, and 4.

So if you follow ex-ante deontology in Case One, you will push at 12:00 but then stop at 13:00. So it is certain that, no matter what, the two people on the track will die and the person on the footbridge will be mildly traumatized. If you simply walked away at 12:00, at least the person on the footbridge would have been better off. Why bother to push the lever at all?

Ex-ante deontology has at least two problems in Case One. The first is that, if you follow ex-ante deontology, you will predictably renege on your earlier decision. This seems irrational.¹⁴ The second problem is that this fickleness

13 If you suspect that ex-ante deontology can avoid this implication if combined with a suitable method of choice for sequential decision problems, hold fire until the next two sections.

14 This is a violation of the principle of dynamic consistency. See McClennen (1990: 120).

can be costly. As we saw, if you follow ex-ante deontology, it is certain that, no matter what, someone will be gratuitously traumatized. This seems immoral.

The second problem, more abstractly, is that ex-ante deontology violates the principle of *Sequential Pareto*, according to which one cannot permissibly follow a course of action that is certain to have the same outcome as some other available course of action, except for someone being worse off.¹⁵ Sequential Pareto should appeal to a deontologist, at least as applied in Case One, since the outcome if you push at 12:00 and stop at 13:00 is certain to be the same, in all respects that a deontologist might care about, as the outcome if you decide not to push at 12:00, except that someone will be worse off. The trolley runs its course either way, for example, and no one is used as a train stop. The conflict with Sequential Pareto is also awkward for ex-ante deontology, which, after all, attempts to respect the idea that one should act in everyone's interest if possible (as per Ex-Ante Pareto). Ex-ante deontology begins to look like an unstable hybrid between the purer extremes of permissivism and ex-post deontology.

I think that Case One poses similar problems for other compromise views. I know two such views, one due to Hare (2013, 2016), and the other due to Setiya (2020). To see what they imply in Opaque and Transparent Footbridge, consider Table 2.¹⁶

TABLE 2 Other compromise views

	Opaque Footbridge		Transparent Footbridge	
	<i>push</i>	\neg <i>push</i>	<i>push</i>	\neg <i>push</i>
Hare's view	✓	✗	✓	✓
Setiya's view	✓	✓	✗	✓

According to either view, the contrast between opaque and transparent versions of Footbridge is less stark than according to ex-ante deontology.

15 A similar principle appears in Gustafsson (2015: 1594–1595). Sequential Pareto is of course stronger than what is needed to argue against ex-ante deontology in Case One. But it is, nonetheless, a plausible principle which also provides a useful reference point. Compare footnote 11. Another problem that ex-ante deontology has in Case One is that, if you follow it, you will follow a course of action whose *prospect* is worse for everyone than the *prospect* of some other available course of action. So ex-ante deontology violates a sequential ex-ante Pareto principle, too. Compare Gustafsson (2018: 599).

16 There might exist further hybrids, similar to ex-ante/ex-post hybrids in the context of egalitarianism. See Bovens (2015). But I fear that they are no more promising than the purer views I consider.

Consider Hare's view first. He thinks that you should push in Opaque Footbridge, but, in Transparent Footbridge, you are permitted to push but also permitted not to push. Hare's explanation is that, in Transparent Footbridge, your reasons to push and your reasons not to push are incommensurable: neither set of reasons is stronger than, weaker than, nor equally as strong as, the other set. Incommensurability arises insofar as you know more about the richly textured lives of the people involved. Hare adds that "faced with such incommensurability, you need a tiebreaker, and a policy of nonintervention in cases like this is as good a tiebreaker as any."¹⁷

I do not think that Hare's view should appeal to deontologists. First, according to Hare, deontic constraints are at best tie-breakers, which is not robust enough for typical deontologists. Second, Hare's explanation of the difference between Opaque and Transparent Footbridge is implausible. In Setiya's words: "it wouldn't affect the ethics of [Transparent] Footbridge if the people involved were perfect duplicates of one another (...) [n]or would it matter if they were people you just met and about whom you know nothing at all."¹⁸ More importantly, however, Hare's view shares ex-ante deontology's problems in Case One. According to Hare's view, you should push at 12:00 but you are permitted not to push at 13:00. But if you push at 12:00 and stop at 13:00, it is certain that, no matter what, someone will be gratuitously traumatized. This is a Sequential Pareto violation.

Consider next Setiya's view. He argues that if you should push in Opaque Footbridge, then you should also push in Transparent Footbridge. Since he thinks that you should *not* push in the latter, he concludes that you are not *required* to push in the former. Instead, he suggests that, in Opaque Footbridge, you are permitted to push but also permitted not to push. Setiya's explanation is that, in Opaque Footbridge, your reasons to push and your reasons not to push are incommensurable. Incommensurability arises because it is difficult to balance a person's interest in survival against the value of his or her autonomy. Setiya adds that "[p]aternalistic intervention is easier to justify when it is the life of a relative or close friend in which you intervene."¹⁹

I think that Setiya's view should appeal to deontologists no more than Hare's. It would not affect the ethics of Opaque Footbridge, it seems, if the people involved were children too young to possess autonomy, nor would it matter if they were your friends and family. More importantly, however, Setiya's view shares ex-ante deontology's problems in Case One. According to Setiya's view, you are permitted to push at 12:00 but required not to push at 13:00. But if you

17 Hare (2016: 467).

18 Setiya (2020: 70). Setiya does not mention that Thomson (1990: 211–220, 185–186) apparently did think that it would matter if the people involved were duplicates. See Kamm (1993: 374–376) for a discussion of Thomson.

19 Setiya (2020: 68).

push at 12:00 and stop at 13:00, it is certain that, no matter what, someone will be gratuitously traumatized. This is another Sequential Pareto violation.²⁰

3 Sophisticated Ex-Ante Deontology

A possible response to ex-ante deontology's problems is to say that, in one's decision, one should take into account what one expects to do in the future. This is the idea behind *sophisticated choice*.²¹ Similarly, according to *sophisticated ex-ante deontology*, one should decide on a course of action whose prospect is best for everyone (if such a course of action exists), *given what one expects to do in the future*, even if that course of action involves doing or intending harm in a way that would otherwise be impermissible.²²

To see how this works in Case One, consider Figure 4. This time the bolded lines indicate moves recommended by sophisticated ex-ante deontology.




FIGURE 4 Case One with Sophisticated Choice

20 According to either Hare's or Setiya's view, in Case One, there is a course of action that you can permissibly follow that will not lead to gratuitous trauma. But, as we saw, in that case, there is another course of action that you can permissibly follow that will. I think that is bad enough. The distinction is similar to that between forcing and non-forcing money pumps. See Gustafsson and Espinoza (2010) and Peterson (2015).

21 It is discussed in McClellan (1990), Rabinowicz (1995), and Buchak (2013: 170–200). Proponents include Strotz (1955) and Hammond (1976).

22 Frick's ex-ante contractualism includes a similar "decomposition test." See Frick (2015: 205).

You know that, no matter what you find out at 13:00, you would then decide not to push the lever, since that would involve doing or intending harm, when it is no longer in everyone's interest. So you know that, if you decide to push at 12:00, you will stop at 13:00. Holding that fixed, your choice at 12:00 is as follows:

Push: Everyone has a $2/3$ chance of death and a $1/3$ chance of mild trauma.

-Push: Everyone has a $2/3$ chance of death and a $1/3$ chance of survival.

Since the latter is in everyone's interest at 12:00, sophisticated ex-ante deontology tells you not to push at 12:00, letting the trolley run its course. So sophisticated ex-ante deontology does not conflict with Sequential Pareto in Case One.

This is cold comfort, however, as it still has to give up Sequential Pareto in the barely different

Case Two. A, B, and C are trapped in three suitcases. The situation is the same as in Opaque Footbridge, except for the following details. It is now 12:00. You still do not know who is in which suitcase but, luckily, everyone will peek out of their suitcases by 13:00. The trolley is slow and will only roll under the footbridge at 13:15. This time the lever is well-oiled but you are not allowed to leave before 13:00. You can either push the lever at 12:00, dropping the footbridge suitcase onto the track, or you can wait to make your decision until 13:00, just before you go. If you do decide to push the lever at 12:00, the two survivors on the track will be mildly traumatized, having to endure an extra hour of agonized screaming from whoever was on the footbridge. That will not happen if you instead decide to push at 13:00.

We can represent this case as in Figure 5.




FIGURE 5 Case Two with Sophisticated Choice

You know that, no matter what you find out at 13:00, you would then decide not to push the lever, since that would involve doing or intending harm, when it is no longer in everyone’s interest. So you know that, if you decide to wait at 12:00, you will decide not to push at 13:00. Holding that fixed, your choice at 12:00 is as follows:

-Push: Everyone has a 2/3 chance of death and a 1/3 chance of survival.

Push: Everyone has a 1/3 chance of death and a 2/3 chance of mild trauma.

If the trauma is mild enough, then, plausibly, it is in everyone’s interest that you push at 12:00. So this is what sophisticated ex-ante deontology tells you to do.

So if you follow sophisticated ex-ante deontology in Case Two, it is certain that, no matter what, the person on the footbridge will be killed and the two people on the track will be mildly traumatized. If you instead decided to wait until 13:00 and push then, at least the two people on the track would have been better off. This shows that sophisticated ex-ante deontology is incompatible with Sequential Pareto, after all. That is implausible. It is also awkward for sophisticated ex-ante deontology which, after all, attempts to respect the idea that you should act in everyone’s interest if possible (as per Ex-Ante Pareto).²³

23 A possible objection is that a course of action that includes pushing at 13:00 is unavailable to a conscientious sophisticated ex-ante deontologist and, so, sophisticated ex-ante

Sophisticated ex-ante deontology also leads to implausible information avoidance in Case Two. Your choice at 12:00 is in effect whether to push immediately or wait to find out who is in which suitcase. Sophisticated ex-ante deontology tells you to do the former, even though it is then certain that, no matter what, two people will be gratuitously traumatized. So if you follow sophisticated ex-ante deontology, you will, in some cases, avoid relevant information at someone's expense. This seems both irrational and immoral.²⁴

Information avoidance may be rational in some cases, for example, in cases where learning is risky, in the sense of potentially making some actual bad things appear unlikely. But why would learning be risky for a deontologist in Case Two?²⁵

Similar problems beset sophisticated versions of Hare's and Setiya's views. Consider Hare's view first. According to its sophisticated version, you are arguably required to push at 12:00. So the view leads to a Sequential Pareto violation and information avoidance. To see why, recall that, according to Hare, you are permitted not to push if you know who is in which suitcase. So you know that, if you wait at 12:00, then you *might* decide *not* to push at 13:00. Holding that fixed, your choice at 12:00 is as follows:²⁶

–*Push*: Everyone *might* have a 2/3 chance of death and a 1/3 chance of survival or everyone *might* have a 1/3 chance of death and a 2/3 chance of survival.

Push: Everyone *definitely* has a 1/3 chance of death and a 2/3 chance of mild trauma.

deontology leads to no Sequential Pareto violation in Case Two. (Compare Seidenfeld (1988: 275–281).) I reply that such a course of action is available to the sophisticated ex-ante deontologist, as it is blocked by no external obstacle. (Compare Steele (2010: 473–4).)

24 The value of a person's prospect should be determined on the basis of the best available information. (Compare Frick (2015: 191–194).) So a possible objection is that there is no good reason to push at 12:00 in Case Two, as that decision would not be based on the information about who is in which suitcase, which will soon be available. I reply that it is important that that information is not yet available at 12:00. At any rate, I could have used a version of Case Two where, arguably, the relevant information does not even exist at 12:00. In that version, you can commit at 12:00 to push at 13:00 or you can leave that decision until 13:00, and an indeterministic mechanism then reshuffles the suitcases between 12:00 and 13:00, rather than before 12:00.

25 See Ahmed & Salow (2019), also Bradley and Steele (2016), and Buchak (2013: 191–200). Good (1967) is the locus classicus.

26 At least on the plausible assumption that you treat everyone equally in the sense that whether you push the lever does not depend on who you see on the footbridge.

Suppose that the people involved are self-interested. And suppose that they know that you are a sophisticated Hare follower. Then they will want you to push at 12:00. If you push at 12:00, that will definitely bring their chance of death to $1/3$, at a small cost of mild trauma if they survive. If you do not push at 12:00, the prospect of that might be better for them but it might also be much worse. So, if you decide not to push at 12:00, the people involved will see that as pointlessly risky. So it is arguably in everyone's interest that you push at 12:00. So that is what a sophisticated Hare follower will do.²⁷

Hare does try to pre-empt objections about information avoidance by saying that, since greater knowledge means greater incommensurability, your better-informed self need not give the right advice to a past self that has to decide when commensurability still obtains.²⁸ This might demonstrate that it makes sense for you to *fail to defer* to your better-informed self. But, in Case Two, a sophisticated follower of Hare will *actively resist* becoming better-informed. That, I think, does not make sense.

Consider next Setiya's view. According to its sophisticated version, you are arguably permitted to push at 12:00. So the view also leads to a Sequential Pareto violation and potential information avoidance. To see why, recall that, according to Setiya, you are required not to push if you know who is in which suitcase. So, if you are a sophisticated follower of Setiya, you know that you would not push at 13:00. Holding that fixed, your choice at 12:00 is as follows:

–*Push*: Everyone has a $2/3$ chance of death and a $1/3$ chance of survival.

Push: Everyone has a $1/3$ chance of death and a $2/3$ chance of mild trauma.

This choice is similar to your choice in Opaque Footbridge where, recall, according to Setiya, you are permitted to push but also permitted not to push. So, presumably, according to a sophisticated version of Setiya's view, you are also permitted to push at 12:00 in Case Two.

²⁷ The people involved will be more insistent if you are known to toss a fair coin to decide between two incommensurable actions. It can then be calculated that, if you push at 12:00, everyone has a $1/3$ chance of death and a $2/3$ chance of mild trauma, but, if you decide to wait at 12:00, everyone has a $1/2$ chance of death and a $1/2$ chance of survival. Plausibly, it is in any person's interest to face the former prospect rather than the latter. Compare Rabinowicz (1995: 595) and Bradley and Steele (2016).

²⁸ Hare (2016: 456–7).

4 Resolute Ex-Ante Deontology

As we saw, ex-ante deontology can tell you to predictably renege on your earlier decisions, even if, as a result, it is then certain that, no matter what, someone will be gratuitously traumatized. A sophisticated version of ex-ante deontology does barely better. It can, in addition, lead to irrational information aversion.

Another possible response to these problems is to say that, in one's decision, one should take into account what happened in the past, rather than what one expects to do in the future. This is the idea behind *resolute choice*.²⁹ Similarly, according to *resolute ex-ante deontology*, one should continue a course of action whose prospect is, or was, best for everyone, even if that course of action, at some point, involves doing or intending harm in a way that would otherwise be impermissible.

To see how this works in Case Two, consider Figure 6. This time the bolded lines indicate moves recommended by resolute ex-ante deontology.




FIGURE 6 Case Two with Resolute Choice

At 12:00 it is in everyone's interest that you wait until 13:00 and push then. This means going up at node 1.

29 Proponents include Machina (1989) and McClennen (1990). See also the discussion in Rabinowicz (1995) and Buchak (2013: 170–200).

At 13:00 it is no longer in everyone's interest that you push. Still, according to resolute ex-ante deontology, you should push, since at 12:00 it *was* in everyone's interest that you wait until 13:00 and push then. This means going up at nodes 2, 3, and 4.

It is easy to see that, unlike ex-ante deontology and sophisticated ex-ante deontology, *resolute* ex-ante deontology does not have problems with Sequential Pareto and information avoidance.³⁰

The view also had a prominent defender: Judith Thomson.³¹ According to Thomson, you should push in Opaque Footbridge, but only as an exception to a blanket prohibition on killing:

“The exceptions (...) are those in which the one who will be killed, and the [two] who will be saved, are members of a group such that it was to the advantage of all the members that the one (whoever he or she would later turn out to be) would later be killed, and the only thing that has since changed is that it is now clear who the one was going to turn out to be.”³²

So, according to Thomson, you should not push the lever in Transparent Footbridge, as there was never a time at which that was in everyone's interest. Sadly, resolute ex-ante deontology also has serious problems.

The first is that resolute ex-ante deontology leads to implausible sensitivity to the distant past. To see this, imagine that it is now 2020 and you face a case like Transparent Footbridge. You can see who is in which suitcase. This is no surprise to you. In 1990 you foretold that in 2020 you will find yourself in a Transparent Footbridge case involving A, B, and C. It would not affect the ethics of the case, it seems, whether you simultaneously foretold who will be in which suitcase. But resolute ex-ante deontology says that it would. That is hard to believe.

³⁰ There are at least two ways to understand resolute choice here. One focuses on the past and its counterfactual continuations into the future, as in Machina (1989); see also Bader (2019). The other focuses on the agent's past plans, as in McClellan (1990). A similar distinction is made in Rabinowicz (1995: 603–604). If you are a resolute chooser of the first type, then, in Case Two, you could say, as you are about to push the lever at 13:00: “If I didn't push, everyone's prospects would have been worse relative to the information available at 12:00.” If you are a resolute chooser of the second type, then you could instead say: “If I didn't push, I would renege on a plan I adopted at 12:00.” I focus on the first, less objectionably inward, way to understand resolute choice.

³¹ See Thomson (1990: 176–202). Her view seems to have changed by the time of Thomson (2008).

³² Thomson (1990: 195).

As this example also shows, in order to establish what is permissible according to resolute *ex-ante* deontology, times in the distant past must be considered. Is it necessary to consider the entire past? If not, how much? This is the second problem.³³

The third and last problem is that resolute *ex-ante* deontology is self-undermining insofar as it undermines intuitions that drive many people towards deontology in the first place. It might no longer be clear, for example, that, in a given situation, it is impermissible to throw a bystander under a trolley to save two other people further down the track. After all, there might have been a time at which it was in everyone's interest that you do so now. So firm intuitions about the impermissibility of doing or intending harm can no longer be considered reliable.³⁴

To end this section, let me note briefly that resolute versions of Hare's and Setiya's views share resolute *ex-ante* deontology's problems. According to a resolute version of Hare's view, you should push if there was once a time at which that was in everyone's interest, but you are merely permitted to do so otherwise. According to a resolute version of Setiya's view, on the other hand, you are permitted to push in the former case but forbidden to do so otherwise. So both views lead to implausible sensitivity to the distant past and, so, to the other two problems as well.

5 Problems for Ex-Post Deontology

So far, I have argued against compromise positions between permissivism and *ex-post* deontology.³⁵ If I am right, the key issue is whether you should push in cases like Opaque Footbridge. If "Yes," then we should accept permissivism,

33 Similar points are made about resolute choice in other contexts by Hammond (1983: 183), Machina (1989: 1651–1653), and Gustafsson (2015: 1599; 2018: 602).

34 Gustafsson (2015: 1599) and Thoma (2019: 249) make a similar point in the context of resolute versions of Scanlon's contractualism and Buchak's risk-weighted expected-utility theory, respectively.

35 Indeed, my argument is related to a more general veil-of-ignorance argument. Compare Harsanyi (1953) and Vickrey (1945). Your choice in Opaque Footbridge is like an impartial spectator's choice in a situation where no individual's final condition is yet known. A persistent objection to veil-of-ignorance arguments is that a hypothetical choice under ignorance is irrelevant to real-life choice under full knowledge. See, for example, Barry (1989: 334–335). It can now be seen that a commitment to sequential Pareto principles and avoidance of implausible sensitivity to the past can link verdicts from behind the veil (as in Opaque Footbridge) to verdicts outside of the veil (as in Transparent Footbridge).

rejecting deontic constraints. If “No,” we are free to accept ex-post deontology instead.

I will now give a new argument against ex-post deontology’s verdict that you should not push in Opaque Footbridge. Let us start with the following case:

Two Opaque Tracks. There are now two parallel tracks. On each track there is an out-of-control lethal trolley, headed for two suitcases. Above each track is a footbridge, on which there is one suitcase. The three suitcases around the first track contain A, B, and C. These suitcases have just been reshuffled, so no one knows who is in which suitcase. The three suitcases around the other track are filled with sand. You have three buttons. Button *a* moves A from his or her suitcase around the first track to the corresponding suitcase around the second track (but you do not see the move). Button *b* likewise moves B, and button *c* likewise moves C. But, if you move any suitcase, that will topple the suitcase that is now on the second footbridge onto the second track. That suitcase, whatever its contents, will then halt the second trolley.³⁶

Two Opaque Tracks decomposes Opaque Footbridge into a number of one-person choices. Each button affects only one person’s position. If you push all buttons, the outcome is the same as if you pushed the lever in Opaque Footbridge. If you push no buttons, the outcome is the same as if you did not push the lever in Opaque Footbridge. Plausibly, then, the verdicts about Two Opaque Tracks and Opaque Footbridge should be in harmony.

So, what should you do in Two Opaque Tracks?

Hare argues that you should push all three buttons, since you should push each button, irrespective of what you do later.³⁷ Take button *a*, for example. If you push it, A has a 1/3 chance of death and a 2/3 chance of survival. If you do not push it, A has a 2/3 chance of death and a 1/3 chance of survival. It is in A’s interest that you push it.

A problem with Hare’s argument is that whether you should push button *a*, say, does seem to depend on what you do later. To see this, suppose that you know that, if you were to push button *a*, you would later push button *b* and button *c*, and that, if you did not push button *a*, you would later push neither button *b* nor button *c*. Holding that fixed, your choice whether to push button *a* is as follows:

Compare Hammond (1996) who offers a sequential version of Harsanyi’s (1955) social aggregation theorem, which is distinct from Harsanyi’s (1953) impartial observer theorem.

36 This is a variation of Hare’s (2016: 461–462) case.

37 See Hare (2016: 457–464). A similar argument can also be found in Frick (2015: 186–194).

Push a: Everyone has a $1/3$ chance of death and a $2/3$ chance of survival.

-Push a: Everyone has a $2/3$ chance of death and a $1/3$ chance of survival.

This choice is like your choice in Opaque Footbridge where, recall, according to ex-post deontology, you should not push the lever. So it cannot simply be assumed, in an argument against ex-post deontology, that, in Two Opaque Tracks, you should push button *a*, *irrespective of what you do later*.

A better argument is available, however. Note first that ex-post deontology leads to *deontic cycling* when combined with the principle of *One-Person Ex-Ante Pareto*, according to which, in a situation where exactly one person is affected, one should choose an option whose prospect is best for that person (if such an option exists).³⁸ To see what I mean, note that, by One-Person Ex-Ante Pareto, you should rather push button *a* than no buttons (since that is best for A and affects neither B nor C), you should rather push buttons *a* and *b* than just *a* (since that is best for B and affects neither A nor C), and you should rather push buttons *a*, *b*, and *c* rather than just *a* and *b* (since that is best for C and affects neither A nor B). Since the choice whether to push all three buttons or none is like your choice in Opaque Footbridge, ex-post deontology implies that you should push no buttons rather than all three. If we write " $x < y$ " to mean that *y* should be chosen when the only alternative is *x*, we obtain the following cyclic pattern of pairwise permissible choices:

-push < push a < push a, b < push a, b, c < -push.

Indeed, Hare's argument is a variation on a well-known type of argument against cyclic preference.³⁹ And the objection I raised against it is a variation on a well-known objection against that type of argument.⁴⁰ Armed with this insight, a more compelling argument against cyclic preference can instead be adapted.⁴¹

³⁸ One-Person Ex-Ante Pareto is similar to Hare's principle of defeasible privacy. See Hare (2016: 463). Deontic cycling is discussed in Willenken (2012), but see also Kamm (1996: 311–354) and Temkin (2012: 194–231).

³⁹ See Rabinowicz (2000b) and Dougherty (2014).

⁴⁰ See Gustafsson and Rabinowicz (2020).

⁴¹ See Gustafsson and Rabinowicz (2020). A similar argument can be found in Broome and Rabinowicz (1999) and Rabinowicz (2000a).

Consider, for example,

The Timer Case. The general setup is the same as in Two Opaque Tracks. But now a special timer controls who goes where around the two tracks. Every half an hour you get a chance to stop the timer. If you stop it at 12:00, no one gets moved to the second track. If you stop it at 12:30, A, B, and C get moved to the second track. If you stop it at 13:00, only A and B get moved to the second track. If you stop it at 13:30, only A gets moved to the second track. 13:30 is the last time to stop the timer. If you do not stop it, no one gets moved to the second track. If you stop the timer at the very start, that will mildly traumatize whoever survives. (Perhaps he or she needs some time to get used to mortality.)⁴²

We can represent this case as in Figure 7.

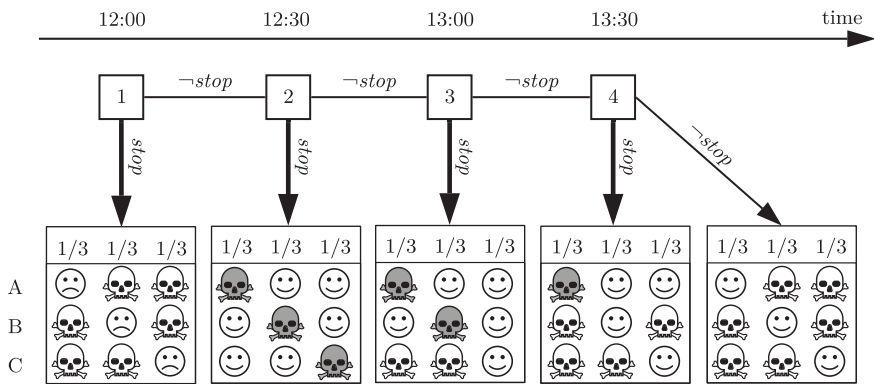


FIGURE 7 The Timer Case

What should you do? You have to stop the timer at 12:00. Any other possibility can be reduced to absurdity, on the assumption of ex-post deontology, One-Person Ex-Ante Pareto, and the plausible premise that you expect to act permissibly in the future, conditionally on having acted permissibly in the past.⁴³

Suppose first that you are permitted to never stop the timer and that you never do. But by 13:30 your only options were:

⁴² The example is similar to the so-called centipede game. See Rosenthal (1981).
⁴³ I do not have the space to discuss this premise in more detail. I think that it is plausible enough to take it for granted here. See Gustafsson and Rabinowicz (2020).

–*Stop*: A has a $\frac{2}{3}$ chance of death and a $\frac{1}{3}$ chance of survival, B and C have a $\frac{2}{3}$ chance of death and a $\frac{1}{3}$ chance of survival.

Stop: A has a $\frac{1}{3}$ chance of death and a $\frac{2}{3}$ chance of survival, B and C have a $\frac{2}{3}$ chance of death and a $\frac{1}{3}$ chance of survival.

It was in A's interest to stop, which would have affected neither B nor C. So, by One-Person Ex-Ante Pareto, that is what you should have done. This contradicts the hypothesis that you are permitted to never stop the timer. So you must stop the timer at 13:30 at the latest.

Suppose next that you are permitted to stop the timer at 13:30 exactly and that you do. You must have acted permissibly up to 13:30, including 13:00. Since you expect to act permissibly in the future, conditionally on having acted permissibly in the past, you expected at 13:00 that you would stop the timer at 13:30 if you did not stop it at 13:00. So you knew at 13:00 that your options were:

–*Stop*: A has a $\frac{1}{3}$ chance of death and a $\frac{2}{3}$ chance of survival, B and C have a $\frac{2}{3}$ chance of death and a $\frac{1}{3}$ chance of survival.

Stop: A and B have a $\frac{1}{3}$ chance of death and a $\frac{2}{3}$ chance of survival, C has a $\frac{2}{3}$ chance of death and a $\frac{1}{3}$ chance of survival.

It was in B's interest to stop, which would have affected neither A nor C. So, by One-Person Ex-Ante Pareto, that is what you should have done. This contradicts the hypothesis that you are permitted to stop the timer at 13:30 exactly. So you must stop the timer at 13:00 at the latest. A similar argument shows that you must stop the timer at 12:30 at the latest.

What about 12:00? Having followed the argument so far, you know at 12:00 that if you do not stop the timer then, you will stop it at 12:30 instead. So you know at 12:00 that your options are:

–*Stop*: A, B, and C have a $\frac{1}{3}$ chance of death and a $\frac{2}{3}$ chance of survival.

Stop: A, B, and C have a $\frac{2}{3}$ chance of death and a $\frac{1}{3}$ chance of mild trauma.

One-Person Ex-Ante Pareto is now silent since more than one person is affected. But this choice is similar to your choice in Opaque Footbridge. So ex-post deontology tells you to take the latter option, stopping the timer at 12:00.

So, if you follow ex-post deontology in the Timer Case, it is certain that, no matter what, the two people on the track will die and the person on the

footbridge will be mildly traumatized. If you simply waited, at least the person on the footbridge would have been better off. This is a violation of Sequential Pareto.

Sequential Pareto is plausible in the Timer Case, as is One-Person Ex-Ante Pareto. Indeed, to contravene One-Person Ex-Ante Pareto in the Timer Case would seem to reveal a form of *causal fetishism*: it would be to care more about impersonal causal structure than about the interest of the only person affected. Even those sympathetic to deontology should find this unpalatable. To see this, compare the Timer Case and a Transparent Footbridge case. If you refuse to kill A to save B and C in a Transparent Footbridge case, you could be suspected of causal fetishism. You can reply, credibly, that you refuse to kill A because you care about A. But if you refuse to take an extra $1/3$ chance to kill A, say, in the Timer Case, you cannot credibly reply that you refuse to do it because you care about A since it is, by hypothesis, in A's interest to take it.

It might seem that a better response to the Timer Case is to combine ex-post deontology with either sophisticated choice or resolute choice. It is easy to see, however, that a sophisticated ex-post deontologist will stop the timer as soon as possible in the Timer Case. On the other hand, a resolute ex-post deontologist, if he or she is to avoid a course of action that leads to gratuitous trauma, would have to contravene One-Person Ex-Ante Pareto at some point. If a resolute ex-post deontologist decides to stop the timer at 13:00, for example, then he or she must then contravene One-Person Ex-Ante Pareto at 12:30.

I conclude that the weakest link is ex-post deontology itself. It should be rejected. Since, in earlier sections, I argued against ex-ante deontology and other compromise views like it, I conclude that permissivism should be accepted. Deontologists will have to resist that argument. But how?

6 Problems for Minimally Paretian Deontology

Even if a deontologist cannot always act in everyone's interest, it might seem that this can only happen when there is uncertainty about the true state of the world. So perhaps a deontologist can still endorse *minimally Paretian deontology*, according to which one should choose an option whose *outcome* is best for everyone (if such an option exists), even if that involves doing or intending harm in a way that would otherwise be impermissible. As it is more limited than ex-ante deontology, it faces none of the problems that I have considered so far. It can also be motivated by the desire to avoid causal fetishism: how could it be wrong to kill someone if that is in everyone's interest, the victim included? Minimally Paretian deontology can be seen as an attempt to

combine deontology with the principle of *Ex-Post Pareto*, according to which one should choose an option whose *outcome* is best for everyone (if such an option exists).

Still, even this form of deontology has a serious problem. We can see it in the following variation of Footbridge:

The Shuffle Case. The setup is the same as in Transparent Footbridge, with two people on the track and one person on the footbridge. The people around the track are reshuffled periodically. First A is on the footbridge, then B, then C, then briefly A again. Every hour on the hour you get a chance to leave, and every hour on the half hour a chance to push a button that topples whoever is then on the footbridge onto the track below. If you ever leave, the shuffle stops. And if you leave right at the very start (when A is first on the footbridge), that will mildly traumatize A. (Perhaps he or she needs some time to get used to mortality.)⁴⁴

We can represent this case as in Figure 8.

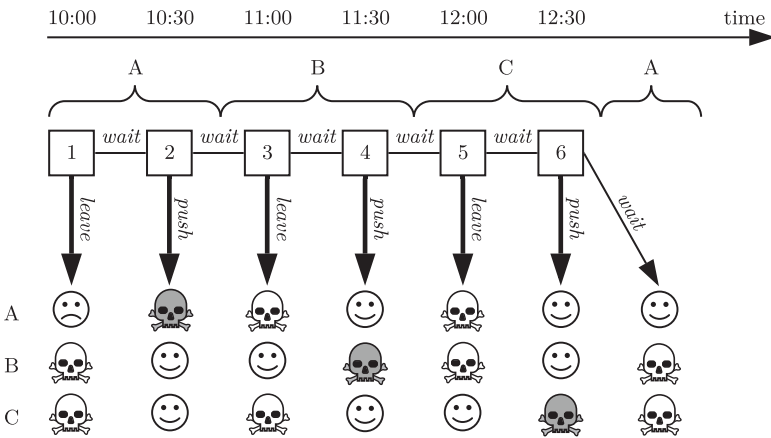


FIGURE 8 The Shuffle Case

44 If you leave, you merely withdraw aid, which is to allow, rather than do, harm. See McMahan (1993).

In this case minimally Paretian deontology leads to deontic cycling. To see this, note that, over time, you have access to the following seven options:

- (a) Wait until A is again on the footbridge. (Allow B and C to die.)
- (b) Push the button when C is on the footbridge. (Kill C as a means.)
- (c) Leave when C is on the footbridge. (Allow A and B to die.)
- (d) Push the button when B is on the footbridge. (Kill B as a means.)
- (e) Leave when B is on the footbridge. (Allow A and C to die.)
- (f) Push the button when A is on the footbridge. (Kill A as a means.)
- (g) Leave when A is first on the footbridge. (Allow B and C to die and allow A to be mildly traumatized.)

And given minimally Paretian deontology, we obtain the following cyclic pattern of pairwise permissible choice:

$$a < b < c < d < e < f < g < a.$$

Imagine, for example, that your choice is between (a) and (b). Then C will die either way, and A will survive either way, but choosing (b) will save B's life. So minimally Paretian deontology recommends (b). Next imagine that your choice is between (b) and (c). (b) is no longer in everyone's interest but, unlike (c), it involves doing or intending harm. So minimally Paretian deontology recommends (c). And so on through the cycle.

A similar argument as in the Timer Case shows that, if you follow minimally Paretian deontology, you have to stop the shuffle at 10:00 in the Shuffle Case.

Suppose first that you are permitted to never stop the shuffle and that you never do. But by 12:30 your only options were:

- Wait:* A survives, B and C die.
- Push:* A and B survive, C dies.

It was in B's interest to push, which would have affected neither A nor C. So, by minimally Paretian deontology, that is what you should have done. This contradicts the hypothesis that you are permitted to never stop the shuffle. So you must stop the shuffle at 12:30 at the latest.

Suppose next that you are permitted to stop the shuffle at 12:30 exactly and suppose that you do. So you must have acted permissibly up to 12:30, including 12:00. Since you expect to act permissibly in the future, conditionally on having acted permissibly in the past, you expected at 12:00 that you would stop the shuffle at 12:30 if you did not stop it at 12:00. So you knew at 12:00 that your options were:

Wait: A and B survive, C dies.

Leave: A and B die, C survives.

You knew that, if you waited, you would later kill C to save A and B and, if you left, you would allow A and B to die and allow C to live. Neither option was in everyone's interest, but the latter did not involve doing or intending harm. So, by minimally Paretian deontology, you should have left. This contradicts the hypothesis that you are permitted to stop the shuffle at 12:30 exactly.

A similar argument shows that you cannot permissibly stop the shuffle at any time later than 10:00. You must stop it at 10:00 by leaving when A is on the footbridge. That means that A will be mildly traumatized. If you simply waited, at least A would have been better off. This is another Sequential Pareto violation.

It is implausible to reject Sequential Pareto in the Timer Case. It is also awkward for minimally Paretian deontologists who, after all, attempt to respect the idea that one should act in everyone's interest if possible, at least provided that there is no uncertainty about the true state of the world. I therefore conclude that deontology cannot be even minimally Paretian.

7 Conclusion

What have I established? I first argued that deontology cannot be combined with the plausible idea that one should act in everyone's interest if possible, even when that involves doing or intending harm. This was the conclusion of my argument against the three forms of ex-ante deontology. I then argued against ex-post deontology. To resist the latter argument, it seems, it would take an unpalatable form of causal fetishism. Lastly, I argued that deontology cannot be combined with the idea that one should act in everyone's interest if possible, even if there is no uncertainty about the true state of the world. It appears that, contrary to popular deontological rhetoric, it is deontology, rather than consequentialism, that does not take individual people seriously.

Acknowledgements

I would like to thank Jessica Fischer, Tomi Francis, Todd Karhu, Aidan Penn, Elliott Thornley, Alice Van't Hoff, Pietro Cibinel, and Adriano Mannino. Special thanks to Ralf Bader, Michal Masny, Andreas Mogensen, Teru Thomas, and Alex Voorhove, and two anonymous referees for this journal for helpful written comments.

Biographical Note

Kacper Kowalczyk is a Research Fellow at University College London, UK

References

- Adler, M. D., & Sanchirico, C. W. "Inequality and Uncertainty: Theory and Legal Applications." *University of Pennsylvania Law Review* 155, no. 2, (2006): 279–377. <https://doi.org/10.2307/40041309>.
- Ahmed, A., & Salow, B. "Don't Look Now." *The British Journal for the Philosophy of Science* 70 no. 2 (2019): 327–350. <https://doi.org/10.1093/bjps/axx047>.
- Alexander, L. "The Means Principle." Legal Studies Research Paper Series, Research Paper No. 14–138 (2014). <https://ssrn.com/abstract=2378608>. Published in abridged form as Alexander, L. "The Means Principle." In F. Kimberley Kessler & M. Stephen J. (eds.), *Legal, Moral, and Metaphysical Truths* (Oxford: Oxford University Press, 2016), pp. 251–264. <https://doi.org/10.1093/acprof:oso/9780198703242.003.0017>.
- Barry, B. *Theories of Justice* (Berkeley: University of California Press, 1989).
- Bader, R. "Agent-Relative Prerogatives and Sub-Optimal Beneficence." *Oxford Studies in Normative Ethics* 9 (2019): 223–250.
- Bovens, L. "Concerns for the Poorly Off in Ordering Risky Prospects." *Economics and Philosophy* 31 no. 3 (2015): 397–429. <https://doi.org/10.1017/S0266267115000188>.
- Broome, J. *Weighing Goods: Equality, Uncertainty and Time* (Oxford: Basil Blackwell, 1991).
- Broome, J., & Rabinowicz, W. "Backwards Induction in the Centipede Game." *Analysis* 59 no. 4 (1999): 237–242. <https://doi.org/10.1093/analysis/59.4.237>.
- Bradley, S., & Steele, K. "Can Free Evidence Be Bad? Value of Information for the Imprecise Probabilist." *Philosophy of Science* 83 no. 1 (2016): 1–28. <https://doi.org/10.1086/684184>.
- Buchak, L. *Risk and Rationality* (Oxford: Oxford University Press, 2013).
- Dougherty, T. "A Deluxe Money Pump." *Thought: A Journal of Philosophy* 3 no. 1 (2014): 21–29. <https://doi.org/10.1002/tht3.91>.
- Fried, B. H. "What Does Matter? The Case for Killing the Trolley Problem (Or Letting It Die)." *Philosophical Quarterly* 62 no. 248 (2012): 505–529. <https://doi.org/10.1111/j.1467-9213.2012.00061.x>.
- Frick, J. "Contractualism and Social Risk." *Philosophy & Public Affairs* 43 no. 3 (2015): 175–223. <https://doi.org/10.1111/papa.12058>.
- Gerbrandy, J. "Identity in Epistemic Semantics." In L. Cavedon, P. Blackburn, N. Braisby, & A. Shimojima (eds.), *Logic, Language and Computation* (Vol. 3) (Stanford, CA: CSLI Publications, 2000).

- Glover, J. *Causing Death and Saving Lives* (Harmondsworth: Penguin, 1977).
- Good, I. J. "On the Principle of Total Evidence." *The British Journal for the Philosophy of Science* 17 no. 4 (1967): 319–321. <https://doi.org/10.1093/bjps/17.4.319>.
- Gustafsson, J. "Sequential Dominance and the Anti-Aggregation Principle." *Philosophical Studies* 172 no. 6 (2015): 1593–1601. <https://doi.org/10.1007/s11098-014-0366-0>.
- Gustafsson, J. E., & Espinoza, N. "Conflicting Reasons in the Small-Improvement Argument." *Philosophical Quarterly* 60 no. 241 (2010): 754–763. <https://doi.org/10.1111/j.1467-9213.2009.648.x>.
- Gustafsson, J. E. "The Difference Principle Would Not Be Chosen Behind the Veil of Ignorance." *Journal of Philosophy* 115 no. 11 (2018): 588–604. <https://doi.org/10.5840/jphil20181151134>.
- Gustafsson, J. E., & Rabinowicz, W. "A Simpler, More Compelling Money Pump with Foresight." *Journal of Philosophy* 117 no. 10 (2020): 578–589. <https://doi.org/10.5840/jphil20201171037>.
- Harsanyi, J. C. "Cardinal Utility in Welfare Economics and in the Theory of Risk-taking." *Journal of Political Economy* 61 no. 5 (1953): 434–435. <https://doi.org/10.1086/257416>.
- Harsanyi, J. C. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility." *Journal of Political Economy* 63 no. 4 (1955): 309–321. <https://doi.org/10.1086/257678>.
- Harris, J. "The Survival Lottery." *Philosophy* 50 no. 191 (1975): 81–87. <https://doi.org/10.1017/S0031819100059118>.
- Hammond, P. J. "Changing Tastes and Coherent Dynamic Choice." *The Review of Economic Studies* 43 no. 1 (1976): 159–173. <https://doi.org/10.2307/2296609>.
- Hammond, P. "Ex-Post Optimality as a Dynamically Consistent Objective for Collective Choice under Uncertainty." In P. K. Pattanaik & M. Salles (eds.), *Social Choice and Welfare* (Elsevier, 1983), pp. 175–205. <https://doi.org/10.1016/B978-0-444-86487-1.50017-4>.
- Hammond, P. "Consequentialist Decision Theory and Utilitarian Ethics." In F. Farina, F. Hahn, & S. Vannucci (eds.), *Ethics, Rationality, and Economic Behaviour* (Oxford: Oxford University Press, 1996), pp. 92–118.
- Hare, C. *The Limits of Kindness* (Oxford University Press, 2013).
- Hare, C. "Should We Wish Well to All?." *Philosophical Review* 125 no. 4 (2016): 451–472. <https://doi.org/10.1215/00318108-3624764>.
- Harman, E. "Review of Caspar Hare, *The Limits of Kindness*." *Ethics* 125 no. 3 (2015): 868–872. <https://doi.org/10.1086/679533>.
- Kagan, S. *The Limits of Morality*. (Oxford: Oxford University Press, 1989).
- Kamm, F. M. "Review: Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status." *Philosophy and Public Affairs* 21 no. 4 (1992): 354–389.
- Kamm, F. M. *Morality, Mortality: Volume II: Rights, Duties, and Status* (Oxford: Oxford University Press, 1996).

- Kamm, F. M. *Intricate Ethics* (Oxford: Oxford University Press, 2007).
- Machina, M. J. "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty." *Journal of Economic Literature* 27 no. 4 (1989): 1622–1668.
- Mahtani, A. "The Ex Ante Pareto Principle." *Journal of Philosophy* 114 no. 6 (2017): 303–323. <https://doi.org/10.5840/jphil2017114622>.
- McClellenn, E. F. *Rationality and Dynamic Choice: Foundational Explorations* (Cambridge: Cambridge University Press, 1990).
- McMahan, J. "Killing, Letting Die, and Withdrawing Aid." *Ethics* 103 no. 2 (1993): 250–279. <https://doi.org/10.1086/293495>.
- Nagel, T. *The View from Nowhere* (Oxford: Oxford University Press, 1986).
- Parfit, D. *On What Matters: Volume One* (Oxford University Press, 2011).
- Peterson, M. "Prospectism and the Weak Money Pump Argument." *Theory and Decision* 78 no. 3 (2015): 451–456. <https://doi.org/10.1007/s11238-014-9435-2>.
- Rabinowicz, W. "A Centipede for Intransitive Preferrers." *Studia Logica*, 67 no. 2 (2000a): 167–178. <https://doi.org/10.1023/A:1010586802757>.
- Rabinowicz, W. "Money Pump with Foresight." In M. Almeida (ed.), *Imperceptible Harms and Benefits* (Dordrecht: Kluwer, 2000b), pp. 123–154.
- Rabinowicz, W. "To Have One's Cake and Eat It, Too: Sequential Choice and Expected-Utility Violations." *The Journal of Philosophy* 92 no. 11 (1995): 586–620. <https://doi.org/10.2307/2941089>.
- Resnik, M. D. *Choices: an Introduction to Decision Theory* (Minneapolis; London: University of Minnesota Press, 1987).
- Rosenthal, R. W. "Games of Perfect Information, Predatory Pricing and the Chain-store Paradox." *Journal of Economic Theory* 25 no. 1 (1981): 92–100. [https://doi.org/10.1016/0022-0531\(81\)90018-1](https://doi.org/10.1016/0022-0531(81)90018-1).
- Scheffler, S. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Revised version of the 1982 edition.) (Oxford: Oxford University Press, 1994).
- Scanlon, T. M. *What We Owe to Each Other* (Cambridge, Mass.; London: Belknap Press of Harvard University Press, 1998).
- Seidenfeld, T. "Decision Theory without "Independence" or without "Ordering": What is the Difference?." *Economics & Philosophy* 4 no. 2 (1988): 267–290.
- Setiya, K. "Ignorance, Beneficence, and Rights." *Journal of Moral Philosophy* 17 no. 1 (2020): 56–74. <https://doi.org/10.1163/17455243-20182841>.
- Singer, P. "Utility and the Survival Lottery." *Philosophy* 52 no. 200 (1977): 218–222. <https://doi.org/10.1017/S0031819100023172>.
- Strotz, R. H. "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies* 23 no. 3 (1955): 165–180. <https://doi.org/10.2307/2295722>.
- Steele, K. S. "What are the Minimal Requirements of Rational Choice? Arguments from the Sequential-Decision Setting." *Theory and Decision* 68 no. 4 (2010): 463–487.

- Temkin, L. S. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012).
- Thomson, J. J. "The Trolley Problem." *The Yale Law Journal* 94 no. 6 (1985): 1395–1415. <https://doi.org/10.2307/796133>.
- Thomson, J. J. *The Realm of Rights* (Cambridge, MA; London: Harvard University Press, 1990).
- Thomson, J. J. "Turning the Trolley." *Philosophy & Public Affairs* 36 no. 4 (2008): 359–374. <https://doi.org/10.1111/j.1088-4963.2008.00144.x>.
- Thoma, J. "Risk Aversion and the Long Run." *Ethics* 129 no. 2 (2019): 230–253. <https://doi.org/10.1086/699256>.
- Vickrey, W. "Measuring Marginal Utility by Reactions to Risk." *Econometrica* 13 no. 4 (1945): 319–333. <https://doi.org/10.2307/1906925>.
- Willenken, T. "Deontic Cycling and the Structure of Commonsense Morality." *Ethics* 122 no. 3 (2012): 545–561. <https://doi.org/10.1086/664750>.