



Opaque Options

Kacper Kowalczyk¹ · Aidan B. Penn²

Accepted: 3 March 2024
© The Author(s) 2024

Abstract

Moral options are permissions to do less than best, impartially speaking. In this paper, we investigate the challenge of reconciling moral options with the ideal of justifiability to each individual. We examine ex-post and ex-ante views of moral options and show how they might conflict with this ideal in single-choice and sequential-choice cases, respectively. We consider some ways of avoiding this conflict in sequential-choice cases, showing that they face significant problems.

Keywords Moral options · Agent-centered prerogative · Deontology · Ex-ante Pareto · Decision theory

Moral options are permissions to do less than best, impartially speaking. They are part of a commonsense, moderate moral outlook, as contrasted with the supposed extremism of consequentialism.¹ This paper argues that—like deontological constraints, contractualism, limited aggregation, prioritarianism, and global risk-aversion²—moral options give rise to a division between ex-ante and ex-post views; the former consider each individual person's expectation, while the latter consider the expected strength of moral considerations within each possible state of the world.³

¹ This contrast is the focus of (Kagan, 1989, pp. 1–46). For a recent defense and overview of moral options, see (Lazar, 2019).

² See, respectively, (Hare, 2016), (Frick, 2015), (Lazar, 2018), (McCarthy, 2006), and (Nebel, 2020).

³ Lazar (2017) endorses an ex-ante account of moral options, but without explicitly addressing the choice between ex-ante and ex-post accounts. For other puzzles for moral options, see, among others, (Pummer, 2016), (Horton, 2017), (Bader, 2019), (Rulli, 2020), and (Muñoz, 2021), which builds on (Kamm, 1985).

Both authors contributed equally to this paper.

✉ Kacper Kowalczyk
kacper.kowalczyk@ucl.ac.uk

Aidan B. Penn
aidan.penn@nyu.edu

¹ Department of Philosophy, University College London, London, UK

² Department of Philosophy, New York University, New York, NY, USA

In Sect. 1, we describe a single-choice situation where ex-post but not ex-ante accounts of moral options appear to conflict with the ideal of justifiability to each individual. In Sect. 2, we describe a sequential-choice situation where ex-ante accounts also seemingly conflict with this ideal and consider some ways of avoiding this conflict, showing that they face significant problems. The difficulty of reconciling moral options with the ideal of justifiability to each individual suggests that a truly individualistic moral outlook must be, in some ways, more demanding than commonsense morality.⁴

1 The opaque options puzzle

Consider

Case One. You and a stranger wash up on two remote islands. You need a dose of a certain medicine to save your leg; the stranger needs a dose of the same medicine to save their life. A nearby ship with exactly one dose is on its way. It's currently headed for the island closest to it. Because of fog, no one knows who's on that island. You have a radio that can be used to send a signal to the ship, redirecting it to the farther island. However, this will delay the ship's arrival, prolonging the suffering of the person saved.⁵

This case can be represented by the decision tree in Fig. 1, in which square choice nodes represent choices you can make (redirect the ship or not), while round chance nodes represent moments when your uncertainty is resolved (when you learn whether you're on the farther or the closer island), with the different possible resolutions having the probabilities shown.

This case raises a puzzle, as there are two plausible but mutually incompatible ways of thinking about it.⁶ On the one hand, you might think that, regardless of which island you're on, you're allowed to redirect the ship. This is because, first, you seem to have the moral option to save your leg instead of saving the stranger's life: saving your leg and saving the stranger's life are each permissible. The sacrifice to you could also be made smaller or larger than that of losing a leg, without affecting the substance of the following discussion. Second, moral options like this are typically stable: you would still be allowed to keep your leg even if that involved some

⁴ This contrasts with Scanlon's optimism about justifying moral options within his brand of contractualism, also motivated by the ideal of justifiability to each individual; see (Scanlon, 1998, pp. 224–225).

⁵ If the ownership of the rescue ship or the medicine is considered important, we could imagine that you own them, or even that the medicine consists of an unowned wooden log floating on the ocean that you can somehow redirect. Thanks to Michael Otsuka for pressing us to clarify this. In other work, one of us uses an analogous case (but without the rescue ship) to examine which theory of moral options is the most plausible (see Penn n.d.), setting aside—what is our central concern here—the tension between moderate theories and justifiability to each individual.

⁶ The puzzle is analogous to Hare's "opaque sweetening problem"; see (Hare, 2010). But the issues are nonetheless different: our puzzle involves interpersonal morality, not intrapersonal prudence; and the issue of preference gaps is not necessarily relevant to the existence of moral options.

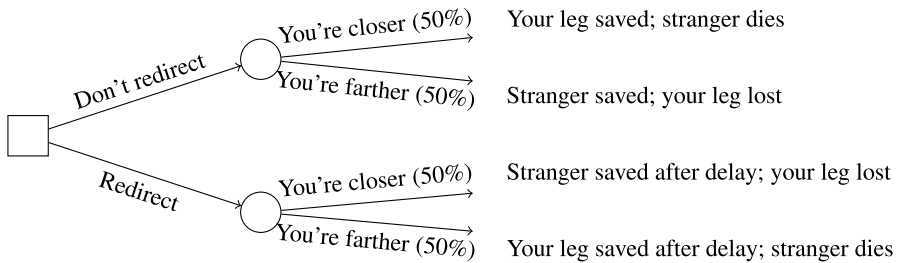


Fig. 1 Case One

additional small cost to you, and you would still be allowed to save the stranger’s life even if that involved some additional small cost to them.⁷ So, if you’re on the closer island, then it’s permissible to redirect the ship because it’s permissible to save the stranger’s life after a delay instead of saving your leg; and if you’re on the farther island, then it’s permissible to redirect the ship because it’s permissible to save your leg after a delay instead of saving the stranger’s life. Surely, then, it’s permissible to redirect the ship. We could support this line of reasoning by appealing to

Statewise Maximality. If, conditionally on every possible resolution of uncertainty about the state of the world, one is allowed to take one option rather than another, then one is allowed to take the first option rather than the second, even before this uncertainty is resolved.

This is a compelling principle. It amounts to a plausible but minimal connection between morally relevant considerations given the actual facts and morally relevant considerations given your uncertainty about the actual facts. If an option is permissible given every way the actual facts could be, then it seems that you are allowed to choose that option, even if you are uncertain about which state of the world is actual.⁸

On the other hand, you might think that you shouldn’t redirect the ship to the farther island. Sending the ship to the farther island gives a 50% chance of saving your leg after a delay and a 50% chance of saving the stranger’s life after a delay; while allowing the ship to reach the closer island gives a 50% chance of saving your leg with no delay and a 50% chance of saving the stranger’s life with no delay. Redirecting is a worse prospect for both of you. Surely, then, you shouldn’t redirect the ship.

⁷ Stability is analogous to “insensitivity to mild sweetening”, what Raz calls the “mark of incommensurability”: two alternatives will remain incommensurable even if one of them is slightly improved; see (Raz, 1985).

⁸ Analogous principles are discussed in (Hare, 2010), (Fleurbaey and Voorhoeve, 2013), (Bales et al., 2014), (Schoenfeld, 2014), (Temkin, 2012, pp. 232–264), and (Francis and Gustafsson, n.d.). For the purposes of this principle, a state of the world must be understood to be fine-grained enough that the state conjoined with information about the selected option settles all morally relevant questions about the resulting situation.

To support the claim that redirecting is a worse prospect for both of you, we could appeal to

The Better-Prize Principle. If one outcome is better for a person than another, then a prospect that gives some chance of getting the first outcome is better for that person than a prospect that gives the same chance of getting the second outcome, other things being equal.

This is another compelling principle. Any reason for favoring the second prospect is more than counterbalanced by a reason to favor the first prospect.⁹ We could then rule out redirecting the ship to the farther island by appealing to

Ex-Ante Pareto. If the prospect of one option is worse for every affected individual than the prospect of another option, then one is not allowed to choose the first option if one can choose the second instead.

This is also a compelling principle. To motivate it, notice that a moral view incompatible with this principle is, in an important sense, not justifiable to each individual, as it sometimes allows us to override what each individual would self-interestedly choose for their own sake.¹⁰

So, we have a puzzle. Proponents of moral options must deny one of the following three seemingly plausible principles: the better-prize principle, ex-ante Pareto, and statewise maximality. To see how they might do this, it's helpful to consider how the following, overly simple, view could be extended to cases of risk:

The Moderate View. An option is permissible if, and only if, no other available option produces a greater weighted total of welfare, relative to some admissible way of weighting the agent's welfare.¹¹

⁹ Compare (Resnik, 1987, pp. 91–92). In this paper, we take prospects to distribute probabilities over states of the world, as in (Arrow, 1966). For our purposes, this principle can be understood to apply only when other things are equal, for example, only to cases which do not involve incommensurability in what's valuable for individuals; compare (Hare, 2010).

¹⁰ A version of this principle was stated by Harsanyi (1955). Compare (Frick, 2013, 2015) and (Hare, 2016). For our purposes, this principle can also be understood to apply only when other things are equal, for example, only to cases which do not involve doing or intending harm; compare (Hare, 2016). It can also be understood to apply only when some people other than the agent are affected, thus allowing for a moral option not to benefit oneself; see, for example, (Slote, 1984). The principle might, nonetheless, appear to beg the question against proponents of moral options, as highlighted by Muñoz's (2019, p. 108) example where you can relieve your headache by taking a pill, thereby also relieving someone else's headache. Muñoz's intuition—shared by Archer (2016) and others—is that not taking the pill is permissible, which contradicts ex-ante Pareto. In response, we observe, first, that some accounts of moral options—like Scheffler's (1994)—lack this implication; second, remaining accounts thereby appear not justifiable to each individual independently of considerations of risk; third, ex-ante Pareto could nonetheless be reconciled with these accounts by adding further provisos, for example, that there is no inconvenience to the agent and that the benefits at stake are substantial. For even if it's plausible to think that you are permitted not to take the headache-curing pill, it's far less plausible to think that you're permitted to take a migraine-inducing pill when not taking the pill is convenient. Thanks to an anonymous referee for pressing us to clarify this.

¹¹ A similar view is formulated by Scheffler (1994, p. 20), but analogous issues arise for other views, for example, those developed by Hurka and Shubert (2012).

The idea here is that there are multiple admissible ways of weighing one's own welfare against that of strangers, which include assigning greater weight to one's own welfare compared to a stranger's. It's permissible for you to save your leg instead of the stranger's life because, relative to some admissible weighting, saving the stranger's life doesn't produce greater weighted total welfare.

There are at least two natural but mutually incompatible ways to extend this view to cases of risk. Put briefly: the ex-ante approach looks at a situation person by person, while the ex-post approach looks at it state by state. Put more precisely: the ex-ante approach determines the scope of our moral options based on the prospects faced by everyone involved, while the ex-post approach determines the scope of our moral options based on what their scope would be in different possible states of the world, together with the probabilities of these states. To better appreciate this distinction, let's examine some examples, starting with the ex-ante approach:

The Ex-Ante Moderate View. An option is permissible if, and only if, no other available option produces a greater weighted total of *expected* welfare, relative to some admissible way of weighting the agent's *expected* welfare.

Given assumptions consonant with our discussion so far, this view implies that you shouldn't redirect the ship. So, this view is compatible with the better-prize principle and ex-ante Pareto but not statewise maximality.

The following view is an example of the ex-post approach:

The Ex-Post Moderate View. An option is permissible if, and only if, no other available option produces a greater *expected* weighted total of welfare, relative to some admissible way of weighting the agent's welfare across different states of the world, not necessarily in the same way in different states.

Given assumptions consonant with our discussion so far, this view implies that you may—but do not have to—redirect the ship. This is because the range of admissible weights includes weights that license saving your leg and weights that license saving the stranger, even if one or the other is delayed; and, so, there is some way of selecting weights across possible states of the world on which redirecting the ship produces a greater expected weighted total of welfare. So, this view is compatible with statewise maximality but not the conjunction of the better-prize principle and ex-ante Pareto.¹²

The ex-ante approach nonetheless seems to have an important advantage over the ex-post approach: namely, the former, but not the latter, seems compatible with the ideal of justifiability to each individual. As we understand it here, this ideal prohibits options that one wouldn't choose if one were concerned solely for each affected individual. In Case One, for example, if one were concerned solely for each affected individual, one wouldn't redirect the ship because doing so gives everyone a worse

¹² It might be suggested that uncertainty is a feature of the outcomes in Case One and, as a result, an ex-post evaluation can be taken to support a requirement not to redirect. But then the resulting view would have analogous implications as the ex-ante moderate view, so it would face problems similar to those described in the next section. Thanks to an anonymous referee for pressing us to clarify this.

prospect. Yet the ex-post approach, unlike the ex-ante approach, allows redirecting in that case.

Proponents of the ex-post approach might argue that this advantage of the ex-ante approach is not decisive.¹³ Instead, we want to argue that, although the ex-ante approach is compatible with the ideal of justifiability to each individual in single-choice cases like Case One, it is not compatible with some compelling ways of understanding this ideal in other cases. There is thus a general tension between this ideal and moral options, regardless of whether they are understood ex-ante or ex-post.

2 Problems for ex-ante views

To see why the ex-ante approach might be incompatible with the ideal of justifiability to each individual, consider

Case Two. This case is like Case One, but your radio is not strong enough to send a signal until the fog clears. Once it clears, you will know whether your island is closer or farther from the ship.

In the earlier case, you choose before you learn; in this case, you learn before you choose. Accordingly, this case can be represented by the decision tree in Fig. 2, which differs from that in Fig. 1 only in that the order of chance and choice nodes has been reversed.

In this case, the ex-ante moderate view implies that you're allowed to redirect the ship regardless of which island you learn that you're on. This is because the ex-ante moderate view extends the moderate view of riskless cases which implies that, under full information, you are allowed to redirect the ship. As a result, however, this view allows you to follow a plan—that is, it allows you to make each choice included in some assignment of choices to reachable choice nodes—that's ex-ante worse for everyone. This is because following the plan of redirecting the ship regardless of what you learn is certain to create a delay for whomever is saved, when compared with following the plan of not redirecting the ship regardless of what you learn.

We could rule out following plans like this by appealing to

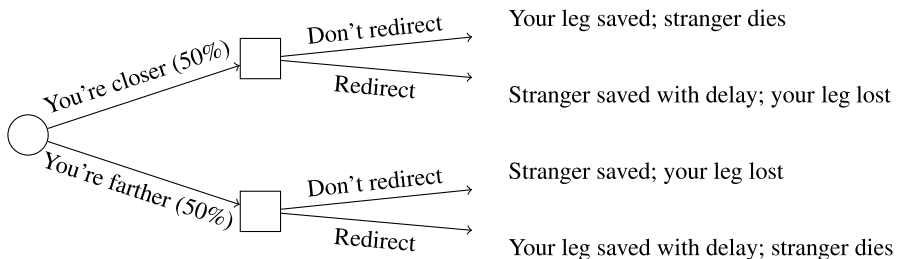


Fig. 2 Case Two

¹³ For further discussion, see (Penn n.d.).

Sequential Ex-Ante Pareto. If the prospect of following one plan is worse for every affected individual than the prospect of following another plan, then one is not allowed to follow the first plan if one can follow the second instead.¹⁴

Note that, since plans can call for a single choice, this sequential principle implies ex-ante Pareto. It also seems compelling for the same sorts of reasons: a moral view incompatible with this principle is, in an important sense, not justifiable to each individual because it sometimes allows one to follow a plan that one would not choose to follow if one were concerned solely for each individual.

The conflict with sequential ex-ante Pareto arises not merely on the ex-ante moderate view as we defined it, but on any account of moral options that satisfies

Decision-Tree Separability. An option is permissible to choose at a node of a decision tree if, and only if, it would be permissible if the sub-tree starting at that node were the whole decision tree.¹⁵

Given the moderate view of riskless cases, this principle implies that, regardless of what you learn, you are allowed to redirect the ship and, so, you are allowed to follow the plan of redirecting the ship regardless of what you learn. This is a compelling principle, especially if we remember that it's compatible with the rest of the decision tree affecting the right way to individuate the available options. For example, depending on the rest of the decision tree, an act might reveal undue favoritism (planning to save a specific individual despite tossing a coin to decide whom to save) or amount to breaking a promise (committing to save one individual and then saving someone else). Neither favoritism nor betrayal are relevant to our example, however.¹⁶

Two possible lines of response thus seem open to the ex-ante moral moderate. First, they might accept sequential ex-ante Pareto but deny decision-tree separability. This would be the implication of endorsing a resolute approach to decision-making, according to which one should act in accordance with a plan favored at some initial privileged moment of time rather than make a choice based on the information now available.¹⁷ The resolute version of the ex-ante moderate view would thus prohibit choosing in a way that, from the viewpoint of that initial moment, would worsen the prospects of everyone affected. One problem with this response is that intuitions that speak in favor of the existence of moral options do not seem sensitive to the past in the way the resolute approach implies. For example, to uphold sequential ex-ante

¹⁴ Compare (Gustafsson, 2022a).

¹⁵ See (Gustafsson 2022a; 2022b, p. 9); but see also (McClennen, 1990, p. 122). This principle is characteristically satisfied by so-called sophisticated accounts of decision-making; see (Hammond, 1976). Two versions of the principle could be distinguished—separability from chance history and separability from choice history—depending on whether the part of the decision tree outside of the relevant sub-tree consists solely of chance nodes or solely of choice nodes; see (Cubitt, 1996). Our discussion of Case Two only needs separability from chance history while our discussion of Case Three below needs full decision-tree separability.

¹⁶ Re-individuation is one way to take into account the value of fairness; see (Broome, 1991, pp. 110–115).

¹⁷ Different types of resolute decision-making are discussed by McClennen (1990) and Machina (1989). A brief survey of some of their problems can be found in (Gustafsson, 2022b, pp. 66–74).

Pareto in the present case, the resolute approach must either prohibit redirecting if you learn you're closer, or prohibit redirecting if you learn you're farther (or both). But these implications seem at odds with the types of intuitions that motivate moral options. Under full information, if you're closer, it seems you are allowed to redirect the ship to save the stranger with a delay instead of saving your leg; the requirement not to redirect would be a requirement to do what's worse overall. Similarly, under full information, if you're farther, it seems you are allowed to redirect the ship to save your leg with a delay instead of saving the stranger; the requirement not to redirect would fail to respect one's integrity, the independence of the personal point of view, or one's autonomy.¹⁸ So, rejecting decision-tree separability will come at a significant cost to the ex-ante moderate.

The second possible response is to accept decision-tree separability and deny sequential ex-ante Pareto—while still accepting ex-ante Pareto. This would be the implication of endorsing a time-slice approach to decision-making, according to which individual choices rather than extended plans are primary objects of moral assessment.¹⁹ One challenge for defending this version of the ex-ante approach is to motivate a version of the ideal of justifiability to each individual that supports ex-ante Pareto but not sequential ex-ante Pareto. In fact, the problem here turns out to be even more difficult. For proponents of the ex-ante moderate approach that accept decision-tree separability must also give up

Sequential Ex-Post Pareto. If, regardless of the true state of the world, the outcome of following one plan is worse for someone than the outcome of following another plan while being equally good for everyone else, then one is not allowed to follow the first plan if one can follow the second instead.²⁰

In this way, the ex-ante moral moderate who accepts decision-tree separability might follow plans that, no matter what, give everyone affected not only a worse prospect (contra sequential ex-ante Pareto) but also a worse outcome (contra sequential ex-post Pareto). To see how the ex-ante moderate view can conflict with sequential ex-post Pareto, consider

Case Three. This case is like Case Two, but your radio's battery is running out. If you use it to play some music, you will later be unable to help the ship navigate, so it will reach the closer island after a minor delay. But if you save the radio for later use, then—after the fog clears—you can also direct the ship to the farther island, causing a longer delay for the person saved.

¹⁸ See, respectively, (Williams, 1981), (Scheffler, 1994), and (Slote, 1985). For an overview of these and other accounts of moral options, see (Lazar, 2019). But see (Bader, 2019) which develops an account of moral options sensitive to the agent's past choices in order to accommodate the impermissibility of sub-optimal beneficence.

¹⁹ Compare (Hedden, 2015) and (Moss, 2015); but also (Buchak, 2013, pp. 170–200). This time-slice approach might, however, be less appealing in the case of interpersonal morality than in the case of intrapersonal epistemology where it has mostly been applied.

²⁰ Compare (Gustafsson, 2015).

In the earlier case, you learn before you choose; in this case, you can avoid learning, at the cost of a minor delay for whomever is saved. This case can be represented by the decision tree in Fig. 3.

The ex-ante moderate view takes into account the expected effects of different available options for the affected people. Given the moderate view of riskless cases, decision-tree separability implies that each option is permissible after the fog clears and, so, if other things are equal—for example, you are not known to favor yourself whenever possible—you should expect yourself to be equally likely to pick any of the permissible options that will then be available.²¹ So, you can expect that saving the radio for later gives

- a 25% chance of saving your leg,
- a 25% chance of saving your leg with a delay,
- a 25% chance of saving the stranger’s life, and
- a 25% chance of saving the stranger’s life with a delay.

This is arguably a worse prospect for both of you than using the radio straightaway, which amounts to

- a 50% chance of saving your leg with a minor delay, and
- a 50% chance of saving the stranger’s life with a minor delay.²²

So, in this case, the ex-ante approach implies that you should use your radio straight-away, thus imposing a cost on one of you, merely so that you avoid having to make

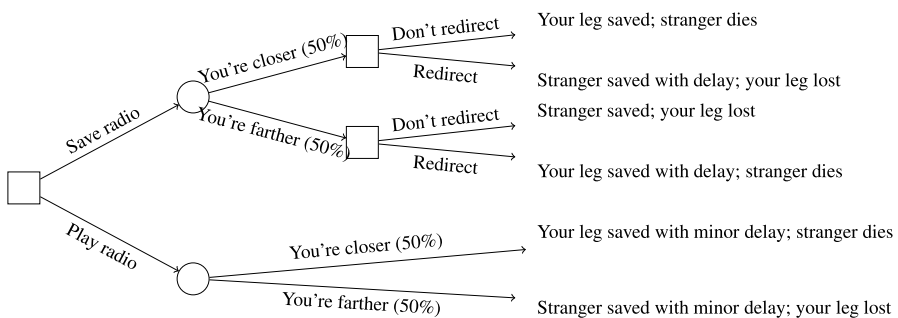


Fig. 3 Case Three

²¹ A similar idea is briefly considered by Rabinowicz (1995). An alternative proposal is that you should assign maximally imprecise probabilities to your future choices between equally permissible options; see (Bradley and Steele, 2016). But then—given prominent permissive decision theories, according to which an option may be chosen if it may be chosen on some way of making the relevant probabilities precise—the option of using the radio straightaway would at least be permissible, again contradicting sequential ex-post Pareto; on imprecise decision theory, see, e.g., (Elga, 2010).

²² This doesn't follow from the better-prize principle alone. But it's highly plausible that there's some way of specifying the effects of the delay and the minor delay such that this claim holds, especially because the former prospect can be made to be worse in terms of expected welfare as well as minimum possible welfare (even though it is better in terms of maximum possible welfare).

a more informed choice about whether to redirect the ship once the fog clears. This violates sequential ex-post Pareto and amounts to a seemingly implausible aversion to information.

So, if proponents of the ex-ante approach accept decision-tree separability, they must reject any conception of justifiability to each individual that supports sequential ex-post Pareto and sequential ex-ante Pareto. As before, they could avoid these implications by denying decision-tree separability, thus also finding a way to condemn costly information avoidance. But, as before, the relevant violations of decision-tree separability are in tension with the central intuitions that motivate moderate moral views in the first place.

3 Conclusion

We have raised problems for reconciling moral options with the ideal of justifiability to each individual that underlies Pareto principles. Ex-post views of moral options conflict with this ideal in single-choice situations, while ex-ante views conflict with it in sequential-choice situations.

Proponents of moral options are thus confronted with a series of difficult questions. The first is whether ex-ante Pareto is correct as applied to scenarios like our Case One. If No, they have to give up on the ideal of justifiability to each individual. If Yes, they confront another difficult question: whether decision-tree separability is correct as applied to scenarios like our Case Two. If No, they must reckon with the problems of resolute choice, revising some of the central intuitions that motivate moderate moral theories in the first place. If Yes, they must give up on the ideal of justifiability to each individual that supports sequential ex-ante Pareto and sequential ex-post Pareto as well as accept costly information avoidance.

In view of these problems, we seem to be left with a choice between a moral outlook that is individualistic yet demanding and one that is moderate yet less centered on the individual. This paper has not provided an argument in favor of the ideal of justifiability to each individual, instead aiming to clarify the implications of accepting this ideal. To this extent, our ambition has been the modest one shared by John C. Harsanyi, a pioneer of the type of arguments discussed in this paper: to clarify a fundamental choice between what he called humanistic and non-humanistic standards of behaviour.²³

Acknowledgements Thanks to Cian Dorr, Tomi Francis, Johan Gustafsson, Kit Fine, Peter Graham, Caspar Hare, Todd Karhu, Chad Lee-Stronach, Richard Roth, Samuel Scheffler, Daniel Viehoff, and Patrick Wu. Special thanks to Michael Otsuka and two anonymous referees for helpful written comments.

Funding Kacper Kowalczyk would like to acknowledge financial support from Longview Philanthropy.

²³ See (Harsanyi, 1958). The puzzle central to this paper is related to a version of Harsanyi's (1955) social aggregation theorem presented by Fleurbaey (2009), but it does not presuppose full expected-utility theory for individuals, completeness of the social preference relation, nor anonymity. Indeed, Harsanyi applied his theorems to social preference only, arguing for moral options as part of his broader rule-utilitarian moral theory; see (Harsanyi, 1977, pp. 652–654).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Archer, A. (2016). Supererogation, sacrifice, and the limits of duty. *Southern Journal of Philosophy*, 54(3), 333–354. <https://doi.org/10.1111/sjp.12176>
- Arrow, K. J. (1966). Exposition of the theory of choice under uncertainty. *Synthese*, 16(3–4), 253–269. <https://doi.org/10.1007/BF00485082>
- Bader, R. (2019). Agent-relative prerogatives and sub-optimal beneficence. In M. Timmons (Ed.), *Oxford studies in normative ethics* (Vol. 9, pp. 223–250). Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198846253.003.0011>
- Bales, A., Cohen, D., & Handfield, T. (2014). Decision theory for agents with incomplete preferences. *Australasian Journal of Philosophy*, 92(3), 453–470. <https://doi.org/10.1080/00048402.2013.843576>
- Bradley, S., & Steele, K. (2016). Can free evidence be bad? Value of information for the imprecise probabilist. *Philosophy of Science*, 83(1), 1–28. <https://doi.org/10.1086/684184>
- Broome, J. (1991). *Weighing goods: Equality, uncertainty and time*. Oxford: Basil Blackwell. <https://doi.org/10.1002/9781119451266>
- Buchak, L. (2013). *Risk and rationality*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199672165.001.0001>
- Cubitt, R. (1996). Rational dynamic choice and expected utility theory. *Oxford Economic Papers*, 48(1), 1–19. <https://doi.org/10.1093/oxfordjournals.oep.a028555>
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, 10. Retrieved from <http://hdl.handle.net/2027/spo.3521354.0010.005>
- Fleurbaey, M. (2009). Two variants of Harsanyi's aggregation theorem. *Economics Letters*, 105(3), 300–302. <https://doi.org/10.1016/j.econlet.2009.08.023>
- Fleurbaey, M., & Voorhoeve, A. (2013). Decide as you would with full information! An argument against ex ante Pareto. In N. Eyal, S. A. Hurst, O. F. Norheim, & D. Wikler (Eds.), *Inequalities in health: Concepts, measures, and ethics* (pp. 113–128). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199931392.003.0009>
- Francis, T., & Gustafsson, J. E. (n.d.). A defence of the principle of stochastic dominance. (Manuscript)
- Frick, J. (2013). Uncertainty and justifiability to each person: Response to Fleurbaey and Voorhoeve. In N. Eyal, S. Hurst, O. Norheim, & D. Wikler (Eds.), *Inequalities in health: Concepts, measures, and ethics* (pp. 129–146). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199931392.003.0010>
- Frick, J. (2015). Contractualism and social risk. *Philosophy and Public Affairs*, 43(3), 175–223. <https://doi.org/10.1111/papa.12058>
- Gustafsson, J. E. (2015). Sequential dominance and the anti-aggregation principle. *Philosophical Studies*, 172(6), 1593–1601. <https://doi.org/10.1007/s11098-014-0366-0>
- Gustafsson, J. E. (2022). Ex-ante prioritarianism violates sequential ex-ante Pareto. *Utilitas*, 34(2), 167–177. <https://doi.org/10.1017/S0953820821000303>
- Gustafsson, J. E. (2022). *Money-Pump Arguments*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108754750>
- Hammond, P. J. (1976). Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1), 159–173. <https://doi.org/10.2307/2296609>
- Hare, C. (2010). Take the sugar. *Analysis*, 70(2), 237–247. <https://doi.org/10.1093/analys/70.2>

- Hare, C. (2016). Should we wish well to all? *Philosophical Review*, 125(4), 451–472. <https://doi.org/10.1215/00318108-3624764>
- Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4), 309–321. <https://doi.org/10.1086/257678>
- Harsanyi, J. C. (1958). Ethics in terms of hypothetical imperatives. *Mind*, 67(267), 305–316. <https://doi.org/10.1093/mind/LXVII.267.305>
- Harsanyi, J. C. (1977). Morality and the theory of rational behavior. *Social Research*, 44(4), 623–656. Retrieved from <https://www.jstor.org/stable/40971169>
- Hedden, B. (2015). Options and diachronic tragedy. *Philosophy and Phenomenological Research*, 90(2), 423–451. <https://doi.org/10.1111/phpr.12048>
- Horton, J. (2017). The all or nothing problem. *Journal of Philosophy*, 114(2), 94–104. <https://doi.org/10.5840/jphil201711427>
- Hurka, T., & Shubert, E. (2012). Permissions to do less than the best: A moving band. In M. Timmons (Ed.), *Oxford studies in normative ethics* (Vol. 2, pp. 1–27). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198790587.001.0001>
- Kagan, S. (1989). *The Limits of Morality*. Oxford: Oxford University Press. <https://doi.org/10.1093/0198239165.001.0001>
- Kamm, F. M. (1985). Supererogation and obligation. *Journal of Philosophy*, 82(3), 118–138. <https://doi.org/10.2307/2026351>
- Lazar, S. (2017). Deontological decision theory and agent-centered options. *Ethics*, 127(3), 579–609. <https://doi.org/10.1086/690069>
- Lazar, S. (2018). Limited aggregation and risk. *Philosophy and Public Affairs*, 46(2), 117–159. <https://doi.org/10.1111/papa.12115>
- Lazar, S. (2019). Moral status and agent-centred options. *Utilitas*, 31(1), 83–105. <https://doi.org/10.1017/s0953820818000201>
- Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4), 1622–1668. Retrieved from <https://www.jstor.org/stable/2727025>
- McCarthy, D. (2006). Utilitarianism and prioritarianism I. *Economics and Philosophy*, 22(3), 335–363. <https://doi.org/10.1017/S0266267106001015>
- McClennen, E. F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511983979>
- Moss, S. (2015). Time-slice epistemology and action under indeterminacy. In T. Szabó-Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 5, pp. 172–94). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198722762.003.0006>
- Muñoz, D. (2019). What we owe to ourselves: Essays on rights and supererogation (Doctoral dissertation, Massachusetts Institute of Technology). Retrieved from <https://hdl.handle.net/1721.1/122426>
- Muñoz, D. (2021). Three paradoxes of supererogation. *Noûs*, 55(3), 699–716. <https://doi.org/10.1111/nous.12326>
- Nebel, J. M. (2020). Rank-weighted utilitarianism and the veil of ignorance. *Ethics*, 131(1), 87–106. <https://doi.org/10.1086/709140>
- Penn, A. B. (n.d.). Prerogatives under risk. (Manuscript)
- Pummer, T. (2016). Whether and where to give. *Philosophy and Public Affairs*, 44(1), 77–95. <https://doi.org/10.1111/papa.12065>
- Rabinowicz, W. (1995). To have one's cake and eat it, too: Sequential choice and expected-utility violations. *Journal of Philosophy*, 92(11), 586–620. <https://doi.org/10.2307/2941089>
- Raz, J. (1985). Value incommensurability: Some preliminaries. *Proceedings of the Aristotelian Society*, 86(1), 117–134. <https://doi.org/10.1093/aristotelian/86.1.117>
- Resnik, M. D. (1987). *Choices: An Introduction to Decision Theory*. London: University of Minnesota Press. Retrieved from <https://www.jstor.org/stable/10.5749/j.ctttshgd>
- Rulli, T. (2020). Conditional obligations. *Social Theory and Practice*, 46(2), 365–390. <https://doi.org/10.5840/soctheorpract20204189>
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge: Belknap Press of Harvard University Press. <https://doi.org/10.2307/j.ctv134vmrn>
- Scheffler, S. (1994). *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Revised version of the 1982 ed.). Oxford: Oxford University Press. <https://doi.org/10.1093/0198235119.001.0001>

- Schoenfield, M. (2014). Decision making in the face of parity. *Philosophical Perspectives*, 28(1), 263–277. <https://doi.org/10.1111/phpe.12044>
- Slote, M. (1984). Morality and self-other asymmetry. *Journal of Philosophy*, 81(4), 179–192. <https://doi.org/10.2307/2026119>
- Slote, M. (1985). *Common-sense morality and consequentialism*. London: Routledge & Kegan Paul. <https://doi.org/10.4324/9781003049265>
- Temkin, L. S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199759446.001.0001>
- Williams, B. (1981). Persons, character and morality. In *Moral luck: Philosophical papers 1973–1980* (pp. 1–19). Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781139165860.002>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.