

# IDENTIFIED PERSON “BIAS” AS DECREASING MARGINAL VALUE OF CHANCES \*

H. Orri Stefánsson

Stockholm University,  
Swedish Collegium for Advanced Study,  
and Institute for Futures Studies

## Abstract

Many philosophers think that we should use a lottery to decide who gets a good to which two persons have an equal claim but which only one person can get. Some philosophers think that we should save identified persons from harm even at the expense of saving a somewhat greater number of statistical persons from the same harm. I defend a principled way of justifying both judgements, namely, by appealing to the decreasing marginal moral value of survival chances. I identify four desiderata that, I contend, any such justification should satisfy, and explain how my account meets these desiderata, unlike some previous accounts.

---

\*I am grateful for the useful comments I received when presenting this paper at the London School of Economics (February 2022), Princeton University (April 2022), Stockholm University (May 2022), Rutgers University (June 2022), Northeastern University (September 2022), Lingnan University (December 2022), University of Oxford (December 2022); and at the Nordic network for bioethics and political philosophy (June 2022). Additional thanks to Erik Angner, Katharina Berndt-Rasmussen, Richard Bradley, Greg Bognar, Nir Eyal, Anna Mahtani, Marcus Pivato, Douglas Portmore, and Joonas Räsänen for fruitful correspondence about the paper. Finally, many thanks to the two reviewers for this journal, for very helpful comments and suggestions. Financial support from Riksbankens Jubileumsfond is gratefully acknowledged.

# 1 Introduction

Suppose that there are two patients who each will die if they don't get a kidney, but there is only one kidney available. Further suppose that there are no good reasons to favour one patient over the other. They are in equal need of the kidney, have an equal claim to it, and, more generally, are equivalent in any morally relevant sense. How should we decide who gets the kidney? A natural thought is that we should use a lottery—for instance, toss a coin—to decide this. Although it is tragic that one patient will have to die as a result of not getting a kidney, the situation is made morally better when whoever dies at least had a chance.<sup>1</sup>

Now consider a decision-problem that may seem very different from the previous one. Suppose that a mine has a collapsed and closed in a single miner who is sure to die if we don't rescue him. Doing so is however quite expensive, and it means that we won't afford to improve safety at the mine. Further suppose that there are 100 miners whose fatality risk would be reduced if we were to spend money on improving safety rather than on rescuing the trapped miner. We estimate that the safety measures would reduce fatality risk for each miner by (almost) one percentage point. So, in expectation, we save (almost) one miner if we improve safety, which is the same as the number of lives that we are (almost) certain to save if we try to rescue the trapped miner. Further suppose that all the 101 miners are equivalent in any morally relevant sense; so, saving each is equally morally valuable. Which should we choose: improve safety at the mine or rescue the trapped miner?

Some think that we have a strong moral reason to save the trapped miner rather than improving safety at the mine (see for instance endnotes 3 and 4). But that may seem a little puzzling. The expected result of not improving safety at the mine is that some not yet trapped miner finds themselves in a situation similar to the currently trapped miner. So why think we should favour the miner who happens to be *now* at a particularly high risk? In other words, why should we prioritise an "identified" (already at high risk) miner over a "statistical" (not yet at high risk) miner (cf. [Schelling, 1968](#))?

The judgement that we should prioritise the already trapped miner can be seen as an instance of the *the identified person (or victim or life) bias*,<sup>2</sup> which Cohen et al. (2015: 1-2) define as:

A greater inclination to assist (and avoid harming) persons and groups identified as those at high risk of great harm than to assist (and avoid harming) persons and groups who will suffer (or already suffer) similar harm but are not identified (as yet).

The general public and policy-makers alike tend to favour saving identified over merely statistical lives.<sup>3</sup> The verdict of philosophers however diverges considerably when it comes to the justifiability of the identified person bias. Some argue that we may have moral reason to prioritise identified persons,<sup>4</sup> while others argue that we should only focus on the expected number of lives saved.<sup>5</sup>

The main aim of this paper is to offer a novel moral justification of the identified person bias. (I will however continue to call it a “bias”, but this should not be understood pejoratively.) A secondary aim of this paper is to show that the question of how to distribute the kidney in the opening example has more in common with the mine example than it might on the face of it seem to have. In particular, I shall argue that both the judgement in favour of using the lottery in the first example and the judgement in favour of prioritising the already trapped miner in the second example can be justified by the *decreasing marginal moral value of survival chances*. However, since a justification of lotteries similar to the one I am proposing here has been discussed elsewhere (Nissan-Rozen, 2019), I will focus mostly on the proposed justification of the identified person bias (which, as far as I am aware, has not previously been discussed).

The rest of the paper is structured as follows. In section 2 I explain the idea of chances having decreasing marginal moral value and show how this idea justifies what I take to be the correct judgements in the examples with which I opened this paper. In section 3 I discuss what kind of probabilities I take the relevant survival chances to be. In section 4 I show how my justification satisfies four desiderata that, I contend, any justification of the identified person bias should satisfy, and argue that some previous attempts fail in this regard. In section 5 I compare my view to both *ex ante* and *ex post* prioritarianism, and show that my view avoids some of the most important objections directed against the other two views, while in section 6 I explain why I think the ethics of distributing chances should be structurally closer to prioritarianism than to egalitarianism. Section 7 discusses what survival chance are morally relevant

in cases where the chances change. In section 8 I consider four objections, and finally in section 9 I conclude.

## 2 The decreasing marginal moral value of chances

Most people recognise a version of decreasing marginal value, namely, decreasing marginal *utility*, that is, value to a person; a phenomenon that seems to hold for many goods (according to most people). For instance, if you own no shoes, then a pair of shoes will presumably be very valuable to you; the utility of a pair of shoes is then very high. In contrast, if you already own three pairs of shoes, then an extra pair will be much less valuable to you than the first. Furthermore, the difference in value to you between, say, the first pair and the eleventh pair will presumably be even greater than the difference in value to you between the first pair and the fourth pair. So, pairs of shoes have decreasing marginal utility.

Quantities of some other goods, however, have decreasing marginal utility within some ranges but increasing marginal utility within other ranges. For instance, £10 000 is typically much more valuable to someone who is starving than it is to a rich person. Still, £10 000 may be more valuable to someone who needs precisely that amount to be able to afford a down payment on their dream house than to someone who can only afford to rent either way. As we shall see in section 8, something similar may be true of the marginal moral value of survival chances: although it is typically decreasing, it may sometimes be increasing, in particular, when a person is close to surviving for sure.

The above examples of shoes and money concerned value *to a person*, not moral value. But many scholars—both philosophers and economists—take *welfare* to have decreasing marginal *moral* value. That is the most common way to operationalise what Parfit (1991) called “the priority view” (and, in fact, the operationalisation he suggested himself). The idea is that given any fixed welfare gain, it is morally better (results in a greater increase in moral value) if the gain is enjoyed by someone who is worse off than if it is enjoyed by someone who is better off. For instance, if Ann is at welfare level 50 and Bob is at level 40, then a 10 unit welfare increase results in greater moral value when Bob gets the increase than when Ann gets it. More generally, the decreasing

marginal moral value of welfare implies that given a fixed sum total of welfare, moral value is higher the more equally that sum is distributed.

The idea that chances can have decreasing marginal moral value has so far hardly received any attention.<sup>6</sup> The general idea is however very similar to the priority view: given any fixed gain in survival chance, it is morally better if the gain is enjoyed by someone whose chance is low than if the gain is enjoyed by someone whose chance is high. For instance, if Ann's chance of surviving the next year is 0.5 while Bob's chance is 0.6, then a 0.1 increase in survival chance results in a greater increase in moral value when Ann gets the increase than when Bob gets it. More generally, the decreasing marginal moral value of survival chances implies that given a fixed sum total of survival chances, an equal distribution of it is morally better than an unequal distribution (other things being equal).

In addition to being structurally similar to the priority view, the view that chances have decreasing marginal moral value can be intuitively justified in a way similar to the priority view. The defender of the priority view can point out that if you can give a fixed welfare benefit to either Ann or Bob, it would seem wrong to give the benefit to the person who is already better off (assuming that Ann and Bob don't differ in any other morally relevant respect). Would it not be fitting to sympathise more with the person who is worse off? Doesn't the worse off person have a stronger claim to our assistance? And doesn't that mean that it would be better to benefit that person rather than the person who is already better off? Similarly, if you can give a fixed increase in survival chance to either Ann or Bob, it would seem wrong to give the benefit to the person who is already less likely to die (assuming that they don't differ in any other morally relevant respect). For again we might ask: Would it not be fitting to sympathise more with the person who is more likely to die? Doesn't that person moreover have a stronger claim? And doesn't that mean that it would be better to increase that person's chances?

In what follows I explain how the decreasing marginal moral value of survival chances justifies both the judgement in favour of using the lottery in the kidney example and the identified person bias in the mine example. However, it may be worth emphasising right away that the justification is meant to leave open how to evaluate (risk-free) *outcomes*.<sup>7</sup> The conclusion argued for below can be combined with different ways to evaluate outcomes

(e.g., utilitarianism, egalitarianism, or prioritarianism). In fact, it can even be made part of a view that says that some consideration(s) take *lexical priority* over improving the survival chances of those whose chances are low. To take a perhaps extreme example: ensuring that someone is saved might take lexical priority over improving people's survival chances, which would imply that we should never give up the option of saving someone for sure for the sake of improving the survival chances of those whose chances are low. So, the below arguments should be read as showing that, *other things being equal*, we should use lotteries to distribute life-saving indivisible goods and we should prioritise identified people over merely statistical people. In other words, the decreasing marginal moral value of chances is *a* reason to distribute goods by lotteries and to prioritise identified people. But other moral considerations may, in some cases, pull the *all things considered* moral judgement in the opposite direction.

Finally, it may be worth emphasising, before considering the below argument, that the identified person bias and the judgement in favour of using the lottery can be justified even if the marginal moral value of survival chances is not *everywhere* decreasing.<sup>8</sup> For instance, these two moral judgements are consistent with thinking that the marginal moral value of survival chances starts *increasing* when someone is very close to surviving for sure. That in fact turns out to have some attractive implications, as we shall see in section 8. So, until we get to section 8, let's leave open the possibility that although at most points on the zero-one interval, the marginal moral value of chances decreases, it need not do so at points sufficiently close to the upper end of the interval.

Let's return then to the examples with which I opened this paper. First consider the example with the two patients and a single kidney. If survival chances have decreasing marginal moral value, then, since everything else is assumed to be equal (in particular, the two patients have an equal claim to and equal need for the kidney), the lottery results in a greater sum of moral value than the options of giving the kidney to either patient without holding a lottery. The decreasing marginal moral value of survival chances implies that the gain in moral value when someone goes from zero survival chance to 0.5 survival chance is greater than the gain in moral value when someone goes from 0.5 survival chance to surviving for sure. Therefore, moral value is higher when each patient gets 0.5 survival chance than when one patient (whoever it is) is sure to die while the other is sure to live. So, the idea that survival

chances have decreasing marginal moral value vindicates the judgement that we should hold a lottery to decide which patient gets the kidney. This is true even if the marginal moral value of survival chances starts increasing when a person's chances are close to 1, as long as that increase is not too extreme. I shall get back to this issue in section 8.

Next consider the mine example. By assumption, the trapped miner is sure to die if we don't rescue him but likely to survive if we do. In contrast, each of the other 100 miners have, we assume, much better survival chances, within the relevant time period, even though we cannot ensure that they survive. Moreover, let's assume that their survival chances are not sufficiently high such that the marginal value of improving their chances is increasing. Then since chances in this range have decreasing marginal moral value—and since for any two miners, it is equally valuable that we save the life of one as that we save the life of the other—the moral value of saving the already trapped miner is greater than the sum of moral value if we instead improve safety at the mine.<sup>9</sup> The total moral value that is gained when someone's survival chance increases from zero to (close to) one is more than 100 times greater than the moral value that is gained when someone whose survival chance is greater than zero increases their chances by (close to) one percentage point. That simply follows from the assumption that survival chances have decreasing marginal moral value (within the relevant range).

Before concluding this section, it might be worth emphasising that my claim is not that most people actually reason in terms of the decreasing marginal moral value of survival chances in examples like those I have been discussing.<sup>10</sup> In fact, it may well be that when people display something that looks like the identified person bias, this is often due to, say, faulty statistical reasoning or the present bias (i.e., the tendency to prioritise present rewards, over future rewards, merely because they are present). Nevertheless, my suggestion is that the identified person bias is in many cases justified since chances typically *do* have decreasing marginal moral value; whether and to what extent people base their judgements on this decreasing marginal value is an empirical question that I shall not try to answer in this paper.

### 3 What are “survival chances”?

Before defending my justification of the identified person bias, let me say a few words about what I mean by “chances”. For the purposes of this paper, I shall mostly be considering *survival* chances, in particular, how a choice affects the probability that someone survives some particular time period. We can equivalently describe this as fatality risks: the probability that someone dies within some particular time period.

For now, I want to remain mostly agnostic about to what extent the arguments in this paper extend to the chances of other goods and bads. The argument plausibly extends to goods and bads that are in some sense sufficiently similar to survival and death; say, avoiding or experiencing severe illness or paralysis.<sup>11</sup> However, some parts of my argument do not seem to extend to all chances. In particular, as we shall see in section 8, there may be reason to think that although survival chances typically have decreasing marginal moral value, such chances have increasing marginal moral value when a person is very close to surviving for sure. The same does not seem as plausible when it comes to, say, chances for welfare. That is, it does not seem plausible, I think, that to maximise moral value we must prioritise someone who is almost certainly going to lead an excellent life over someone who is unlikely to lead an excellent life, simply because the first is so likely to lead an excellent life. But if my argument does not extend to chances for welfare, then it cannot extend to the chances of all goods and bads that affect a person’s welfare. I am not in a position to delineate with any precision the goods and bads to which my argument fully extends. Therefore, I shall continue to write about survival chances; but this can be understood more generally as the chances of goods and bads that are relevantly similar to survival and death.

But what kind of probabilities are these survival chances? The probabilities I have in mind should ideally satisfy two conditions that some may find hard to reconcile. The first condition, which in fact the probabilities in question *must* satisfy, is that they should be compatible with the claim that for each person in the mine example, the probability that the person survives the relevant period could be strictly between zero and one. Similarly, it must be possible that for each patient in the kidney example, the probability that they receive the kidney if we let a coin toss decide is (at least close to) 0.5. If this condition were not



satisfied for “survival chances” then the idea of decreasing marginal moral value of survival chances would neither justify the identified person bias nor would it justify letting a coin toss decide in the kidney example.

The second condition, which the probabilities in question should ideally satisfy, is that it should make sense to claim that it can be good *for a person* when their survival chances increase. For instance, that each patient has a chance of getting the kidney is not only morally good; it is also good for each patient. By this I do not only mean that it can be *ex ante* good for a person—that is, good for them *in expectation*—to have a chance at getting some good outcome.<sup>12</sup> Most would agree with that. Instead, I mean that it can be good for a person to have a chance at getting some good outcome even if that outcome does not materialise.<sup>13</sup> One reason for imposing this condition will become apparent in section 5. As we shall see, however, neither my justification of using the lottery nor my justification of the identified person bias *depends* on this condition being satisfied, although the justification will arguably be more attractive when it is satisfied.

Some may find it hard to see that these two conditions could be satisfied by the same kind of probability. Coin tosses are governed by deterministic processes, and perhaps it is predetermined when the mine collapses and who (if anyone) will then become fatally injured. That means that according to some metaphysical views (e.g., Schaffer, 2007), the *objective* probability that each miner survives any period is either zero or one; similarly, the objective probability that a patient wins the kidney in a coin toss is either zero or one. So, if this view on the metaphysics of probabilities is correct, then the first of the above conditions would seem to imply that the probabilities of interest are what some would call *epistemic* (e.g., Otsuka, 2015b), that is, the probabilities that are supported by the available evidence. Or they might simply be the subjective probabilities—that is, the degrees of belief—of the relevant decision-maker. In any case, the probabilities in question cannot be objective given a metaphysical view like, say, Schaffer’s (2007).

Now, it may not be at all obvious that epistemic (and perhaps even subjective) probabilities are inferior to more objective probabilities when it comes to capturing the intuition that we should give both patients a chance in the example with which I opened this article.<sup>14</sup> Perhaps it would be just as good to let a coin toss that has already taken place—but whose outcome is yet unknown—

decide which patient gets the kidney, even though doing so would not give the two patients an equal objective chance of surviving. There are, however, at least two problems with using non-objective probabilities in the view I want to defend. One problem is that it would imply an arguably problematic aversion to information. I shall get back to this problem in section 8.

Another problem with using non-objective probabilities in this context is that it would be hard to reconcile with the second condition that I introduced above. That is, it seems hard to make sense of the idea that it could be good for a person to have a merely subjective or epistemic probability of getting some good outcome. To see this, suppose first that the relevant survival chances are epistemic, that is, relative to the available evidence. Imagine that some new piece of evidence becomes available and suggests that Daphne's survival chance is better than what the previously available evidence suggested. Or, if you think that such evidence concerns groups rather than individuals: Suppose that new evidence becomes available and shows that philosophers are more likely to survive until eighty than the previously available evidence suggested.<sup>15</sup> Is it really plausible that, in the first case, the availability of the evidence is *good for* Daphne? She would surely welcome the news. But, objectively, it would seem that what happens isn't good—nor bad—for her, except in the sense that it may make her *feel* good. The world hasn't become better for her; rather, the world is and was all along better for her than we thought. And similarly in the example with the survival chances of philosophers. Even more clearly: the fact that some decision-maker revises upwards their subjective probabilities for Daphne's survival is not objectively good for her (again, setting aside the possibility that it may make her feel good).

So are these two conditions on probabilities mutually compatible? Can we find a respectable notion of probability that is objective while being consistent with non-trivial probabilities in deterministic systems? I think we can. In fact, these two conditions are satisfied by, for instance, the theory of "Humean chances" developed by Roman Frigg and Carl Hoefer (see, e.g., 2010; 2015; see also Hoefer, 2007), building on the work of David Lewis (e.g., in 1980).<sup>16</sup> It would take me too far from the topic of this paper to describe this Humean view in detail, let alone to try to defend it. But an important aspect of the view, which might be worth mentioning to explain how it satisfies both of the above conditions, is the idea that the objective chances are those numbers assigned

to events by the probability rules that are part of the system that strikes the best balance between simplicity, strength, and fit to the non-modal facts of our world. Another feature of Humean chances that is important for the present purposes, is that, unlike relative frequencies, they are not defined relative to a (more or less arbitrary) reference class.<sup>17</sup>

Importantly, Frigg and Hoefer argue that even though a coin toss, say, is governed by a deterministic process, the best system will contain a rule that says, roughly, that “where such-and-such symmetry is to be found in a solid object of middling size with  $n$  possible faces that can land upward (or downward, thinking of tetrahedra), and when such objects are thrown/rolled, the chance of each distinct face being the one that lands up (or down) is exactly  $1/n$ . Given what we know about dice and tetrahedra and so forth, it is quite plausible that this rule belongs in the Best System; and it entails the coin-flip chances” (2010: 359). So, the rule will imply that when a normal coin is tossed in a normal way, the chance with which it lands up (and down) is 0.5 (assuming that we condition on the macroscopic state of the universe at the time of the toss). Moreover, since simplicity, strength, and fit are meant to be *objective* criteria—and so is the rule for determine when these criteria strike the best balance—the chance in question is though to be objective.<sup>18</sup>

In what follows, I shall assume that the survival chances of interest are Humean chances, in Frigg and Hoefer’s (but not Lewis’) sense.<sup>19</sup> But what is important is not the complete Humean theory; what is important is just that the chances in question are consistent with, first, non-trivial chances in systems governed by deterministic processes, and, second, the claim that having a chance is (objectively) good for a person (and not just because it makes them feel good).<sup>20</sup>

## 4 Four desiderata and other accounts

Let’s now consider four desiderata that, I contend, any justification of the identified person bias should satisfy.

The first desideratum is that the justification of the identified person bias should not imply that one should prioritise identified persons *no matter* how many lives would otherwise be saved in expectation. This desideratum ap-

appears to be violated by Frick's (2015a; 2015b) justification. Frick argues that when deciding whether to prioritise the already trapped miner or instead increase safety for the other 100 miners, we should do a pair-wise comparison of the strength of the claims that each of the miners have on our assistance, which in this case corresponds to how much is at stake for each miner *ex ante*. Moreover, since Frick is sceptical of the aggregation of claims (in cases like these), he concludes that we should rescue the already trapped miner rather than improving safety at the mine.

This justification of the identified person bias however has seemingly awkward implications for situations where one can save *many* more lives, in expectation, by reducing risk for the statistical lives. For instance, suppose that we face a choice between either saving an already trapped miner or instead increasing safety at the mine and thereby reducing fatality risk for 10 *million* miners by 10 percentage points each. In terms of the expected number of lives saved, the difference between the two alternatives is one to one million. Surely in that case we should save the statistical lives. But that is not what the competing claim model implies (as Frick acknowledges): the trapped miner's claim is still stronger than the claim of each of the other ten million miners, and their claims are not supposed to be aggregated. As a result, Frick ends up endorsing a pluralist view, according to which the competing claims model provides one of various possibly conflicting reasons, which implies that we should *sometimes* prioritise identified people even if it means saving fewer lives in expectation, but not for *any* number of expected lives saved. However, Frick offers no principled way of determining how to weigh what he calls "reasons of fairness", which favour the identified miner, against what he calls "well-being-given reasons" (2015b: 196), which favour saving a greater number of statistical lives. The same is not true of my view, as I explain in a moment.

Now, a committed pluralist might respond that the aforementioned indeterminacy in Frick's view captures an important truth about morality, since it simply *is* indeterminate how to weight reasons of fairness against well-being-given reasons. Moreover, it is worth acknowledging that resorting to pluralism is not an *ad hoc* attempt at saving Frick's view from counterexamples, since pluralism may seem attractive independently of the need to avoid these extreme implications. Still, it seems unfortunate to me that we would have to appeal to different principles in the two miners examples that *only* differ in terms of how

many lives are at stake; that is, given Frick's account, we appeal to reasons of fairness to get the intuitively right verdict in the less extreme miners example but have to appeal to well-being-given reasons to get the intuitively right verdict in the more extreme miners example. In contrast, my account delivers the intuitively right verdict in both examples by appealing to the decreasing marginal moral value of survival chances (as I shall explain shortly). While this may not be seen as a conclusive argument in favour of my approach over Frick's, I hope that it will at least be seen as *a* benefit of my approach.

The second desideratum is that a justification of the identified person bias should also justify the bias in cases where prioritising saving statistical lives leads to *somewhat* more lives saved in expectation. This desideratum is violated by Daniels (2015) (and possibly by Hare, 2012). He points out that in cases where there is a conflict between saving identified lives and saving merely statistical lives, there often is a sense in which the identified persons are worse off, by which he must mean that they are worse off in terms of chances. For instance, when comparing the miner who is already trapped with the 100 miners who face some risk whenever they enter the mine, there is a clear sense in which the already trapped miner is worse off than the other 100; he will die for sure if we don't come to his rescue, whereas the others have a much greater chance of surviving their work in the mine.

However, since Daniels does not offer a way in which one could compare the moral value of helping those whose survival chances are very low with the moral value of saving statistical lives, he notes that: "Risk concentration may matter morally in breaking ties but not do more" (2015: 119). In other words, on Daniels view, it may be that in a choice between helping the trapped miner and reducing fatality risk for the 100 miners by one percentage point each, we should help the already trapped miner; but if we have the option of reducing fatality risk for the 100 miners by, say, 1.0001 percentage point each, we should do so rather than helping the already trapped miner.

In contrast to both Daniels' and Frick's justifications, my justification of the identified person bias implies a simple precisification (see next endnote) of the trade-off between on the one hand what Frick calls "fairness" and on the other hand the number of lives saved, in a way that satisfies both the first and the second desideratum, that is, in a way that avoids the extremes of both Frick's and Daniels' justifications. Since we are assuming that survival chances'

marginal moral value is typically *decreasing* rather than merely non-increasing we get a result that is stronger than Daniels'; that is, we should sometimes save identified lives at the expense of saving in expectation a greater number of statistical lives. But we should not *always* do so. In fact, the feature of my view that justifies the identified person bias, namely, the extent to which the marginal value of chances decreases, also ensures that we avoid Frick's extreme implication. Any reasonable extent to which the marginal moral value of survival chances might decrease implies that we should save the trapped miner rather than improving safety for 100 miners by 1.0001 percentage point each, but that we should not save the trapped miner rather than improving safety for 10 million miners by 10 percentage points each.<sup>21</sup> So, we don't have to appeal to different principles to get the intuitively right verdict in on the one hand the example where the trade-off is between saving one miner vs. making a tiny improvement in safety for 100 miners, and on the other hand in the example where the trade-off is between saving one miner vs. making a significant improvement in safety for 10 million miners.

The third desideratum is that the justification of the identified person bias should be weaker in situations where even the "statistical persons" are at high risk of dying soon, compared to situations where mortality risks are generally low. More concretely, the reason to prioritise identified persons should be weaker in poor countries where, say, the healthcare systems are severely underfunded and annual fatality risks are high, compared to affluent countries with better funded health-care systems. In the latter case, it might be unfair—to the poorly off majority—if we spent as high a proportion of our health care budget on, say, late stage cancer patients as we can justify spending on such patients in countries where annual fatality risks are lower. For instance, [Cohen et al. \(2015\)](#) ask:

Many nations fund expensive intensive care and dialysis for people who otherwise are at high risk of dying soon. They do so even though many preventative interventions would avert more deaths per dollar, often from the same diseases. Is that a mistake? *In particular, should poor and middle-income countries never fund expensive dialysis so long as they lack the funds for more cost-effective interventions for most everyone in need, including medication for high blood pres-*

sure (a cause of kidney failure and the need for dialysis)? (2015: 2, emphasis added)

The decreasing marginal moral value of survival chances implies that we should be willing to sacrifice less in terms of the expected number of lives saved for the sake of saving identified lives in situations where people generally are at a higher risk of dying (within a specified time period) compared to situations where people in general face lower such risks. The reason for this is that where average annual fatality risks are high, the “statistical persons” have low survival chances; that is, those persons who are not identified as having a specially high risk also have low survival chances (compared to countries where annual fatality risks are generally low). Therefore, the gain in moral value when the survival chances of these statistical persons are improved is greater than when the chances are improved for statistical persons that have high survival chances (assuming that their chances are typically not so high as to have increasing marginal value). So, in a country where many people are close to starving, or where health is generally very poor, the decreasing marginal moral value of survival chances justifies sacrificing fewer expected statistical lives for the sake of saving identified lives, compared to countries where few people are close to starving and health is generally better.

The fourth and final desideratum is that an account that seeks to justify the identified person bias must satisfy what Brock (2015: 43) calls “The Principle of the Equal Moral Worth of All Human Lives”:

This [principle] implies that all human persons deserve equal moral concern and respect, as it is sometimes put, and that, all else equal, saving more lives rather than fewer is morally better [...] It also implies that all else equal, identified and statistical lives have equal moral value. [...] Given my assumption about the equal moral value of human lives, a defense of the different moral importance of saving identified versus statistical lives requires an argument why this difference matters that is morally compatible with humans’ equal moral value. (ibid, 43-44)

Some attempts at justifying the identified person bias don’t seem to satisfy the above principle. One example is Slote’s justification, based on his “virtue



ethics that anchors morality in empathic caring reactions” (2015: 150). “Empathy is more sensitive to what is contemporaneous and definite than to what is future and indefinite,” Slote (2015: 153) points out, which, he thinks, not only explains but even justifies our willingness to spend more on saving identified persons than on saving merely statistical lives. But it is unclear whether this is consistent with all lives’ equal moral value. It seems not only inevitable but even fitting that we have stronger “empathic caring reactions” to some than to others. For instance, it would seem fitting that we empathise more strongly with those we love than with those we don’t know. Furthermore, some people may unfortunately not enlist emphatic caring reactions from anyone. So, it is unclear whether Slote’s justification satisfies Brock’s principle.

Does my justification of the identified person bias satisfy Brock’s principle? I think it does. Most would agree, I think, that Parfit’s priority view is consistent with the principle that every persons’ welfare is equally morally important. Similarly, those who criticise the view never do so, as far as I can tell, on the grounds that the view denies that humans have equal moral value. For instance, even though the priority view says that if Ann is at welfare level 10 while Bob is at level 5 then it is better to improve Bob’s welfare by one unit than improving Ann’s welfare by one unit, the view still implies that, say, the fifth unit in Ann’s welfare is equally morally valuable as the fifth unit in Bob’s welfare. So, although we prioritise Bob over Ann when Ann is better off, that does not mean that we fail to show the same concern and respect for Ann as for Bob, since we would have prioritised her had she been worse off than Bob.

Similar observations can be made about my justification of the identified person bias. The fact that we prioritise the trapped miner over the 100 not yet trapped miners, say, in the sense that we spend resources on rescuing him rather than on improving safety at the mine, does not mean that we fail to show equal moral concern and respect for all miners. For all of the 101 miners, it is equally morally valuable to increase their survival chances from, say, 0.1 to 0.2. And the same is true for any other chance-increase from a fixed level. It is just that whenever someone is badly off in terms of survival chance, any increase in their survival chance typically results in greater moral value than when someone who is well off in terms of survival chance sees their chance increase by the same magnitude. But that is not because the *life* of the one who has low survival chance has greater moral worth than the life of the one



who has higher survival chance. Their lives have equal moral worth; that is perfectly consistent with saying that moral value is increased more when we increase the survival chance of whoever is otherwise at a higher risk of dying.

## 5 Comparison with prioritarianism

The reader might wonder how my proposal differs from the version of the priority view that is known as *ex ante* prioritarianism. An *ex ante* prioritarian takes *expected welfare* to have decreasing marginal moral value, and then (typically) applies expected value theory to evaluate risky social prospects. This implies both the judgement in favour of the lottery and the identified person bias (assuming, for instance, that the people whose lives are at stake will have equal lifetime welfare if they survive).

So, how does the justification I am proposing differ from *ex ante* prioritarianism? First, while these two approaches agree in these two cases, the underlying reasons are quite different. An *ex ante* prioritarian thinks that in general, expected welfare has decreasing marginal moral value; hence, as an implication, one should use the lottery in the kidney example and one should prioritise identified over statistical persons. In contrast, the view that I am proposing doesn't say that expected welfare has decreasing marginal moral value. Instead, it simply says that for *some* "goods", such a survival, chances for those goods typically have decreasing marginal moral value.

Second, and relatedly, the approach that I am suggesting allows that the structure of the moral value of chances may differ depending on what good or bad is at stake. For instance, my approach is consistent with saying that when it comes to chances for income, say, one should evaluate them in a utilitarian fashion, even though we should agree with the recommendations of (*ex ante*) prioritarianism as far as life-and-death prospects are concerned.<sup>22</sup> This flexibility seems particularly important in light of the fact that there may be good reasons (which we shall consider in section 8) for thinking that although the marginal moral value of survival chances typically decreases, it increases when a person's survival chance is very close to 1. For it seems less plausible that the marginal moral value of, say, the chance of someone having a high income increases as the person is close to having a high income for sure.

Third, and perhaps most importantly, my proposal avoids what is arguable the most serious problem with *ex ante* prioritarianism, namely, that it sometimes recommends choosing contrary to what we know we would choose—as prioritariness or egalitarianism—if we had more information (see, e.g., Fleurbaey and Voorhoeve, 2013, Otsuka, 2015a, Otsuka and Voorhoeve, 2018).<sup>23</sup> The reason my proposal avoids this problem is that, as I illustrate below, it shares some structural similarities with *ex post* prioritarianism as opposed to *ex ante* prioritarianism.

	State 1, p=0.5	State 2, p=0.5
A	Ann: 39, Bob: 60	Ann: 60 Bob: 39
B	Ann: 30, Bob: 70	Ann: 70, Bob: 30

Table 1: Anti-correlated two person gambles over welfare

Consider, for instance, alternatives A and B in table 1. Both alternatives are 50:50 gambles between Ann being better off amongst the two and Bob being better off. For instance, if we choose A and state 1 obtains, then Ann will be at welfare level 39 while Bob is at welfare level 60. The difference between the two alternatives is that, no matter how the uncertainty is resolved—that is, irrespective of whether state 1 or state 2 obtains—alternative A results in more equality and is considerably better for the worse-off than B, and only results in a slightly lower aggregate welfare than B. Therefore, (most) *ex post* egalitarians and *ex post* prioritariness would prefer A to B. However, B is better than A *ex ante* for both Ann and Bob: both of them receive slightly higher expected welfare from B than from A. Therefore, an *ex ante* prioritarian is committed to preferring B to A (in virtue of being committed to the *ex ante* Pareto principle), which might seem counterintuitive to those who have egalitarian or prioritarian intuitions.

The view that I am proposing does not say that expected welfare has decreasing marginal moral value. Hence, it need not recommend B over A in table 1. However, one might worry that it makes an analogous recommendation about survival chances, for instance in the choice problem represented by table 2, which should be interpreted as table 1 except that the numbers now represent survival chances.<sup>24</sup>

Does the view that I am defending necessarily favour D over C, just like *ex ante* prioritarianism necessarily favours B over A? It does not. Since the view is that chances have decreasing marginal moral value—rather than *ex-*

	Event 1, p=0.5	Event 2, p=0.5
C	Ann: 0.39, Bob: 0.60	Ann: 0.60 Bob: 0.39
D	Ann: 0.30, Bob: 0.70	Ann: 0.70, Bob: 0.30

Table 2: Anti-correlated two person gambles over survival chances

*pected* chances having decreasing marginal moral value—it may well favour C over D, for the same reason that *ex post* prioritarianism may favour A over B. In fact, given any reasonable extent to which the marginal moral value of survival chances might decrease, my view does favour C over D.<sup>25</sup> So, the view that I am suggesting, although having some of the same implications as *ex ante* prioritarianism, is in some ways structurally more similar to *ex post* prioritarianism.<sup>26</sup>

The structural similarity to *ex post* prioritarianism might make the reader worry that, like *ex post* prioritarianism, my view cannot account for the arguably important difference between the unity of the individual and the separateness of persons (see, e.g., Otsuka and Voorhoeve, 2009, Voorhoeve and Fleurbaey, 2012).<sup>27</sup> Consider for instance the alternatives in table 3, each of which is a 50:50 gamble that only affects Ann’s welfare. And suppose that, in situations of risk, an *ex post* prioritarian applies expected value theory, and therefore maximise expected priority weighted welfare (see, e.g., Rabinowicz, 2002 and Adler, 2012). Then the *ex post* prioritarian may favour E over F, even when choosing F would be in the interest of the only person affected by the choice.<sup>28</sup> Moreover, they will evaluate the trade-off between, on the one hand, Ann getting higher *minimum* welfare, and, on the other hand, Ann getting higher *expected* welfare, in precisely the same way as they would evaluate a structurally similar choice where the trade-off is between on the one hand ensuring higher minimum welfare for Ann when she is worst off, and, on the other hand, Ann and Bob getting a higher sum of welfare. In other words, the view is not sensitive to the fact that Ann is unified in a sense that, say, Ann+Bob are not.

	State 1, p=0.5	State 2, p=0.5
E	Ann: 39	Ann: 60
F	Ann: 30	Ann: 70

Table 3: One person gambles over welfare

Does my view face an analogous problem? Consider the alternatives in

	Event 1, p=0.5	Event 2, p=0.5
G	Ann: 0.39	Ann: 0.60
H	Ann: 0.30	Ann: 0.70

Table 4: One person gambles over survival chance

table 4, which is like table 3 except that the numbers should now be interpreted as survival chances for Ann. The decreasing marginal moral value of chances might mean—in fact, for any reasonable extent to which the marginal moral value of chances decreases, it *does* mean<sup>29</sup>—that G has higher expected moral value than H. However, note that in this case, a choice of G over H need not go against Ann’s interest. It may well be more important *for Ann* that her survival chance increases from 0.30 to 0.39 than that her survival chance increases from 0.60 to 0.70. That is, the increase in value *to Ann* may be greater in the first case than in the second. (Note that this is not a claim about Ann’s preferences; her preferences might not perfectly track what is good for her.) The analogous move cannot, of course, be made by the *ex post* prioritarian to justify the choice of E over F, since, by assumption, the numbers in table 3 completely represent Ann’s interest. More generally, since a measure of a person’s welfare is meant to account for anything that makes the person’s life go better or worse, an *ex post* prioritarian cannot consistently claim that the decreasing marginal moral value of welfare corresponds to what is good for the person. In contrast, I can consistently claim that the decreasing marginal moral value of survival chances corresponds to how survival chances are good for a person. I shall get back to this issue at the end of section 8.

Still, some may find it odd that I treat an *intrapersonal* trade-off, such as that in table 4, analogously to how I would treat an *interpersonal* trade-off, such as a choice between on the one hand giving Ann a 0.7 chance while Bob gets a chance of 0.3 and on the other hand giving Ann a 0.6 chance while Bob gets a chance of 0.39. For it might seem that the “difference between the unity of the individual and the separateness of persons requires that there be a shift in the moral weight that we accord to changes in [survival chances] when we move from making intrapersonal tradeoffs to making interpersonal tradeoffs”, to slightly paraphrase Voorhoeve and Fleurbaey (2012: 381). As they note, such a shift cannot be accounted for by a view such as prioritarianism, according to which the moral significance we attach to improving the situation of those who

are badly off does not depend on others being better off. However, it can be accounted for by a view that is *egalitarian*, to use Parfit's (1991) terminology. According to such a view, we should prioritise those who are badly off if and because they are *worse off than others*; or, in the present context, we should prioritise those whose survival chances are low if and because they are lower than other people's chances. The next section considers such a view.

## 6 Why not egalitarianism about chances?

The following argument might seem to support egalitarianism about chances.<sup>30</sup> The reason we should hold the lottery, in the example with the two patients, is that it would be *unfair* to give the kidney to one patient without giving the other patient a chance. Similarly, it would be unfair to improve safety at the mine, thus giving the not yet trapped miners an even greater survival chance, when the already trapped miner's chances are so slim. Such fairness considerations are *comparative*.<sup>31</sup> Therefore, rather than being prioritarians about chances, in the sense that we take chances to have decreasing marginal moral value, we should be egalitarians about chances, in the sense that we judge the moral value of improving a person's survival chance by considering how their chance compares to those of others.

Can we say anything positive in favour of chance prioritarianism over chance egalitarianism? One thing that may count in favour of the former view over the latter, is that the latter would seem to be vulnerable to a version of the *levelling down objection*. That is, the latter view seems to imply that equalising people's survival chances by lowering everyone's chances would be in at least one respect an improvement; namely, with respect to the value of equality. The idea that chances have decreasing marginal moral value does not imply any special value of equality, over and above the value of improving the chances of those who are badly off in terms of chances, and is therefore not vulnerable to a similar objection.

Now, some (welfare-)egalitarians respond to the levelling down objection by pointing out that even if equalising a distribution by levelling down does make the distribution better in one respect (i.e., with respect to equality), it does not follow that levelling down would be an improvement all-things-considered

(see, e.g., [Temkin, 1993](#)). In contrast, others have developed egalitarian theories that do not face the levelling down objection (e.g., [Otsuka and Voorhoeve, 2009](#)). A chance egalitarian can of course adopt either of these strategies too. But in either case, as long as chance egalitarianism is not extensionally equivalent to some version of chance prioritarianism, the former would, unlike the latter, imply that the value of improving someone's chances cannot be evaluated separately from other people's chances, which seems to me to count in favour of the latter view over the former. For instance, in a case like that described in [table 2](#), a chance prioritarian can judge that C is better than D without having to consider the chances of people other than Ann and Bob. Similarly, in the case illustrated by [table 4](#), a chance prioritarian need not consider anyone's chances other than Ann's when comparing alternative G with H. In contrast, a chance egalitarian must in both cases consider the chances of all those who belong to whatever population (or society) the egalitarian takes to be the object of their distributive concern. This is of course a practical drawback, since it makes policy evaluation so much harder. But it would also seem to be a normative drawback. After all, why should, say, Carl's survival chance affect whether we choose G or H in a case like that in [table 4](#) that only affects Ann?

Another disadvantage of chance egalitarianism, compared to chance prioritarianism, is that the perhaps most natural version of the former would not satisfy the third desideratum I introduced above. If an egalitarian thinks that the moral value of improving a person's survival chance depends on how it compares to the chances of other people within the same society, then they would be willing to sacrifice as many statistical lives for the sake of saving identified lives in societies where most people are at high risk of dying soon as in societies where fatality risks are lower. As we saw in [section 4](#), the same is not true of chance prioritarianism.

Finally, and perhaps most importantly, as we saw in the last section, chance prioritarianism is less vulnerable to what I think are the most important objections that egalitarians (such as [Fleurbaey, Otsuka, and Voorhoeve](#)) have raised against welfare prioritarianism. In sum, it seems to me that while chance prioritarianism avoids at least some of the problems of welfare prioritarianism, chance egalitarianism faces all the well-known problems of welfare egalitarianism. Nevertheless, I acknowledge that a potential weakness of my view is that, just like *ex post* prioritarianism, it cannot account for the shift in the moral

weighing that some think should occur when we move from an intrapersonal trade-off to an interpersonal one.

## 7 What happens when chances change?

The idea that survival chances are goods with decreasing marginal value may seem to result in a dilemma. Consider a variant on the original mine example, where this time a lottery is held to decide which miner enters the mine at each point in time. So, for each point in time, all the miners have the same chance of finding themselves in the mine at that time. Now suppose that at time  $t$ , a particular miner happens to find himself in the mine when it collapses and closes him inside; again, we suppose that he is sure to die if we don't rescue him, but if we do rescue him then we cannot afford to improve safety at the mine. In other words, this is exactly like the mine example discussed on page 2, except that this time there is a sense in which survival chances have already been equalised: all the miners had the same chance of finding themselves in the mine at time  $t$ . Do we in this case have equally strong reason—and the same kind of reason—to help the trapped miner as in the original mine example?

I think we do. Although the miners' chances were equal at some point, this is no longer the case once the mine collapses and traps a miner. The world has now changed in a way that makes their chances unequal; the trapped miner is sure to die, unless we rescue him, whereas the others have much better chances. So, the assumption that chances are goods with decreasing marginal moral value implies that in this case too we should prioritise the trapped miner.

But now someone might worry that this has strange implications for the example with two patients and one kidney. Suppose that the decision has been made that a lottery will be held next Thursday to determine which patient receives the implant. Thursday comes, the lottery is held, and patient 1 wins. Right before the implant, the chief physician however realises that although the patients' chances were equal on Wednesday, that is no longer the case; now patient 2 is sure to die whereas patient 1 has very good chances. Should she hold a new lottery to equalise chances again?<sup>32</sup>

The obvious answer is that she should not. But is that consistent with the judgement that in the last version of the mine example we should prioritise the



trapped miner? There we said that things had changed such that the miners' chances were no longer equal. But isn't that true in the kidney example too? The outcome of the lottery that was held on Thursday changed things such that the patients no longer have an equal chance. So what explains why that change doesn't require re-equalising the chances, while the change in the mine example does require such a re-equalisation?

It may be tempting to answer that if we re-equalise chances in the kidney example, by holding a lottery again on say Friday, then we are giving the patients different *total* chances. For if we decide to hold a second lottery, as a result of patient 1 winning the first lottery, then we are in fact giving patient 2 a  $3/4$  total survival chance, whereas patient 1 only gets a total survival chance of  $1/4$ . However, suppose that the second lottery would have been held even if patient 2 had won the first lottery. In that case the two patients have an equal total survival chance. Still, we should clearly not hold the second lottery.

A better answer, I think, is that the circumstances that determine the patient's claims have not changed in the kidney example. I suggest that the right way to understand an idea such as chance prioritarianism, is that when people have competing claims to our contributions to their survival chances, we should distribute the available resources in a way that maximises moral value in light of the hypothesis of decreasing marginal moral value of chances.<sup>33,34</sup> Similarly, ordinary (*ex post*) prioritarianism should be understood as the view that when people have competing claims to societies' resources, we should distribute them in a way that maximises moral value in light of the hypothesis of decreasing marginal moral value of welfare (see, e.g., [Adler, 2012](#)).

When the mine collapses and closes in a single miner, that changes things in a way that is relevant to how we should adjudicate between on the one hand the claims of the trapped miner who needs rescue and on the other hand the rest of the miners who benefit from improved safety at the mine. The world has changed in a way that affects the strength of people's claims to our assistance. Therefore, the hypothesis of the decreasing marginal moral value of survival chances should be (re-)applied, and now it implies (given the other assumptions of the example) that we should prioritise the trapped miner. In contrast, the fact that the lottery has resulted in an outcome which means that one patient dies for sure does not change things in a way that is relevant to how we should adjudicate between the claims of the two patients. The world



has *not* changed in a way that affects the strength of people's claims. What has changed is merely that a value maximising process—indeed, the process that satisfies the competing claims as well as possible—has taken place. Therefore, we should not hold the lottery again.

What should we do then if, before the transplant takes place, we learn that a third patient needs the kidney and has an equal claim to it as the other two (cf. [Vong, 2015](#))? I suggest that we then use a weighted lottery, that gives this third patient a  $1/3$  chance while giving patient 1, who won the first lottery, a  $2/3$  chance. For then as long as we would have done the same had patient 2 won the first lottery, each patient gets a  $1/3$  total chance at the kidney. So, moral value is then maximised, given the hypothesis of decreasing marginal moral value of survival chances.

## 8 Four objections

Before concluding, I will address four objections. The first objection is that my account implies that we should sometimes be willing to accept a cost to prevent a fatality even though we know for sure that someone else will die instead. This objection is nicely illustrated by a hypothetical example due to Nir Eyal.<sup>35</sup> He uses the example to criticise the idea that *epistemic* survival chances should be equalised. But the example can also be used to illustrate the cost of my proposal, that is, the idea that we should be concerned with the distribution of objective (Humean) survival chances.

[I]magine that the Battle of Britain is on, and that you are staffing a tall power station in the center of a densely populated British city. A Blitz raid begins. As Luftwaffe planes start dropping shells and buildings catch fire, you look out of your window and notice a streetlamp that was erroneously left undimmed. A nearby house and its residents are therefore at concentrated risk for mortal attacks, because German pilots used lights as convenient targets. You control the power system and could turn off the streetlight. There would be considerable financial cost, but the house residents' lives might be saved, and you judge the cost to be well worth the high chance of saving lives. Then, however, you have a second thought:

What would turning the light off accomplish? German bombs are already pounding your densely populated city (which has no empty lots). If this house is spared, another house will not be. Flicking the lights off for the sake of that house's residents will only reallocate objective danger. It will not reduce it. Is it really something that you have strong moral reason to do, reason enough to justify considerable financial cost? [...] In fact, the [moral value of survival chances] would likewise increase if turning this streetlight off [randomly] turned on another one in front of another house in town [...] a moment later; and if turning that light off would a moment later turn on another, and so forth. (Eyal 2015: 102-103, slightly modified to fit my account)

My account implies that switching off the light is something that you have a moral reason to do even at some financial cost, assuming that it is *indeterminate* which house will be bombed if you do so (for instance, if it is truly *random* which streetlight will be turned on instead). After all, by turning off the light, you then give all the residents the same objective chance of surviving.<sup>36</sup> Some may nevertheless find this implication of my view to be counterintuitive. Why should we be willing to pay a cost that only reallocates, but does not reduce, objective danger? Now, part of that counterintuitiveness may be due to the fact that, by turning off the light, one is arguably actively playing a role in harming whoever happens to be bombed as a result, which is not true if one leaves the light on. If this completely explains the counterintuitiveness of the implication in question, then that need not be troubling for my account. After all, as discussed on page 6, my account is not meant to be a complete moral theory. My claim is just that, *other things being equal*, one should improve the survival chances of those whose chances are low. But that is consistent with thinking that one should also be concerned with avoiding that *oneself* inflicts harm on others.

I suspect that many readers will not feel that the above completely explains the counterintuitiveness of the implication of my account that Eyal's example illustrates. Perhaps another fact that helps explain this counterintuitiveness, without undermining my account, is that in Eyal's example, when we reduce the risk faced by one person (or group of people), we *immediately* expose

another person (or group of people) to the exact same risk. That might explain why the implication in question may seem counterintuitive even to some who think that we should prioritise the trapped miner. But then the intuition would seem mistaken. Other things being equal, it should, I contend, make no moral difference, to the mine example, if one of the 100 not yet trapped miners is for some reason *immediately* put at risk if we decide to rescue the trapped miner.<sup>37</sup> A risk of harm is just as bad whether it occurs at time  $t$  or a moment later—assuming that everything except the timing of the risk is held fixed.

Now, I admit that the above two observations might fail to convince those who find the implication under consideration to be counterintuitive. So, eventually, I will presumably have to bite the bullet, and admit that there are some somewhat counterintuitive implications of my view that I suggest we learn to live with. But I hope that those who are convinced by my argument that we have a moral reason to prioritise the survival chances of those whose chances are low will not find the implication in question too hard to live with. By turning off the light you raise the survival chance of a group of residents who otherwise would have very low chances. Those who are convinced that we have a moral reason to prioritise those who are badly off in terms of welfare, say, are generally willing to do so at a considerable financial cost, and, in fact, even at the expense of aggregate welfare. Similarly, if we think that we have a moral reason to prioritise those whose survival chances are low, then we should be willing to do so even at some cost, and even when that does not increase aggregate survival chances.<sup>38</sup>

The second objection I will consider is that there is a sense in which my proposal implies that one should be willing to pay to avoid getting information.<sup>39</sup> For instance, suppose that two vaccines are equally effective at eliminating some fatal pediatric disease; further suppose that each vaccine has to be given to every child in the population for it to be effective against the disease. Unfortunately, one vaccine kills an identified child who is known to have a very rare gene mutation; the other vaccine is expected to kill one child but it is indeterminate which child will die. In all other respects the vaccines are predicted to have identical effects. Moreover, to keep things simple, let's suppose that all children have equal and high survival chances when we set aside the risk from, first, the disease against which we are choosing a vaccine, and, second, the potential side effects of the vaccine.

My account implies that the second vaccine is morally better than the first; therefore, there is some price one should be willing to pay to get to use the second vaccine rather than the first. In a sense, that means that my account implies that one should be willing to pay a price to not know which child will die. But note that one is paying for more than *just* not knowing this. One is also paying for each child facing an equal *and high* survival chance rather than one child facing a very low (in fact, zero) survival chance. It seems clear to me that there is some price worth paying for that.

But now suppose that after we have chosen the second vaccine, an oracle that can somehow predict the outcome of random processes offers to tell us which child will die because of the vaccine. Should we accept the information? Perhaps we should not trust our intuitions in cases like these, since it may be hard for us to make sense of the outcome of a random process being predictable. I shall get back to that issue shortly. But first, note that someone who thinks that we should be concerned with purely epistemic (or even more subjective) survival probabilities would turn down the oracle's offer. In fact, there is even some price that such a person would be willing to pay for the oracle *not* to give them the information. In contrast, if we are concerned with objective (Humean) survival chances, as I have been arguing we should be, then there could be good reasons for accepting the information—and we would have much less reason to pay to avoid getting the information.<sup>40</sup> After all, it would still be true that we gave each child the same objective survival chance. Moreover, we might use the information to somehow compensate the child who is going to die, for instance, by trying to make the rest of their life as good as possible.<sup>41</sup>

Some may still object that my account implies that we would have a reason to pay a higher price for the second vaccine even if (whether we like it or not) the oracle tells us *beforehand* which child that vaccine will kill. And that may seem counterintuitive.<sup>42</sup> Now, that intuition may be partly explained by the fact that it would be so much more psychologically difficult to administer a vaccine when we *know* whom it will kill (recall endnote 40). So, one benefit of the second vaccine would be lost. In addition, I suspect that this intuition is partly explained by the fact that it is hard to intuitively understand how it could be true that each child had the same objective chance of surviving yet we knew beforehand which child would die. In other words, I suspect that the intuition that we should not, in this case, be willing to pay more for

the second vaccine, can be at least partly explained by the fact that when we assume that we can know beforehand who will die, we implicitly think of the survival probabilities as being epistemic. But I acknowledge that some will likely disagree, in which case I may again have to bite the bullet, and accept this as another counterintuitive implication of my view, which is hopefully offset by the view's other virtues.

The third objection is quite different from the first two.<sup>43</sup> Imagine that you face the following choice. You can either (a) increase the survival chance of one person from 0 to (close to) 1, or (b) increase the survival chance of million people from zero to one in a million. My view implies that you have a moral reason to choose the latter. After all, the decreasing marginal moral value of survival chances implies that one million times the moral value of increasing one person's survival chance from zero to one in a million is greater than the moral value of increasing someone's survival chance from zero to (close to) one. This may seem counterintuitive.

There may be two worries about the above implication of my view. First, one might wonder if this is really consistent with the identified person bias. If we choose (a), then we know whom we will (almost certainly) save. If we choose (b), however, then (if the chances in question are independent) there is roughly 63% chance that we will save someone from a group of one million people but we have no idea whom. So, is choosing (b) really consistent with the identified person bias? Second, one might wonder if it is morally defensible to increase the survival chances of people by such tiny amounts (from 0 to 0.000001) when one has the option of almost ensuring that someone survives instead of facing certain death.

Now, in one sense the first of the above worries is not particularly interesting from a philosophical point of view, since it is partly just terminological: what do we mean by "identified person bias"? Still, the worry is important, given that I claim to have justified the identified person bias. If the above implication of my view is inconsistent with the identified person bias, then I will have to weaken my claim. But is the above inconsistent with the identified person bias? Recall the definition of the identified person bias from above, which Cohen et al. state in their introduction to the volume *Identified vs. Statistical Lives*, and which I reproduce here for ease of reference:

*The identified person bias:* A greater inclination to assist (and avoid harming) persons and groups identified as those at high risk of great harm than to assist (and avoid harming) persons and groups who will suffer (or already suffer) similar harm but are not identified (as yet). (Cohen et al., 2015, 1-2)

Given this authoritative definition, the bias is neutral in a case like the one currently under consideration, that is, in a case where we have to choose between assisting a single person who is otherwise sure to die or one million people all of whom are otherwise sure to die. For according to this definition, the “identified” in the trade-off between identified vs. statistical lives refers not to whom we will save, but rather to those “at high risk of great harm” (ibid.). Thus, since in the case under consideration both the single person and the group of one million people are at high risk of great harm, the identified person bias has no implication for this decision.<sup>44</sup>

Still, one might of course wonder if it is morally defensible to improve the survival chances of the one million people by such a tiny amount, to such an extremely low survival chance level, rather than saving someone almost for sure. A similar worry of course arises for prioritarianism; in fact, it arises for any aggregative moral view. For instance, one might wonder if it is morally defensible to improve the welfare of one million people by 0.0001 unit rather than improving the welfare of one person by one hundred units, as prioritarianism implies we should if all 1,000,001 people are initially equally well off. (Similarly, one might wonder if it is morally defensible to improve the welfare of one million people by 0.0001 unit rather than improving the welfare of one person by ninety-nine units, as utilitarianism implies we should.) Now, I think one could argue that *if* one is (sufficiently) attracted to prioritarianism (or utilitarianism), then to reject favouring the one million over the single person in cases like these is to commit the fourth of Parfit’s five mistakes in moral mathematics, namely “to ignore *very small* [...] effects on very large numbers of people” (Parfit, 1984, 75). The same, I contend, holds if one is attracted to the idea that, first, survival chances have decreasing marginal moral value and, second, that moral value should be maximised, but nevertheless rejects favouring the one million over the single person in the above alleged counterexample to my view. Still, I acknowledge that some may find this to be a counterin-

tuitive implication of my view, in which case one could for instance combine the decreasing marginal moral value of chances with some merely-partially aggregative view (recall discussion on page 6).

The fourth and final objection that I shall consider is that my view has counterintuitive implications when faced with the choice between, on the one hand, saving one person who is otherwise sure die, and, on the other hand, saving a greater number of other people, all of whom are otherwise sure to die.<sup>45</sup> Suppose for instance that a boat with five people has capsized, and four of them are in one lifeboat while one person is in another, and sadly we only have time to reach one of the lifeboats before they capsize too. The decreasing marginal moral value of survival chances then implies that moral value is maximised by using a weighted lottery to decide which group to save, rather than simply saving the greater number (as before, I am assuming here that saving each of the five is equally morally valuable). But that may seem counterintuitive. Shouldn't we make sure we save the greater number of lives?

Before I respond to this objection, it should be noted that even though we assume that chances have decreasing marginal moral value, we maximise moral value in the above example by giving a *much* greater survival chance to the four people compared to the one. Recall (from endnote 21) that to say that survival chances have decreasing marginal moral value is equivalent to saying that there is some concave function on chance values that represents their moral value. The square root function is a simple concave function that is often used for illustrative purposes. Applying that function to this example, we find that moral value is maximised when we use a lottery that gives roughly 0.94 chance to saving the four and only 0.06 chance to saving the one.<sup>46</sup> The reason is that in this example we don't have a *fixed* sum of survival chances to distribute. Instead, we can ensure a much higher sum of survival chances by distributing them less equally, that is, by giving greater chances to the four compared to the one. To take a more extreme example: If one person is in one lifeboat while 99 people are in another, then the square root function implies that moral value is maximised by giving the 99 a roughly 0.9999 survival chance while giving the single person only a 0.0001 chance.

Some may still find the above result counterintuitive. In the more extreme example, by maximising moral value we run a 0.01% risk of saving only one person when we could have saved 99 people for sure. In response, one might



be tempted to argue that the moral value of chances is indeed maximised by holding a lottery that runs a 0.01% risk of saving only one person, but to point out that the value of chances is not the sole determinant of moral value. So, perhaps the moral value of saving 99 people for sure outweighs the moral value of chances in this case.<sup>47</sup> Although in many ways appealing, this response may not be available to me, given my criticism (in section 4) of Frick's treatment of the identified person bias. For this response would suggest that I too have to appeal to different principles depending on how many lives are at stake.

Instead, I would like to suggest, in response to this fourth objection, that although the moral value of survival chances is decreasing at *most* points, it is perhaps not decreasing everywhere. Suppose the moral value of improving someone's survival chance from, say, 0.99 to 1 is greater than the moral value of improving someone's survival chance from, say, 0.59 to 0.6. That is consistent with the stipulation that up to a survival chance of 0.99, the marginal moral value of survival chances decreases. But it will reduce the risk we should be willing to run of saving 1 person rather than saving 99 for sure. It will imply that we should tolerate *some* such risk, in some cases, unless we stipulate that improving someone's survival chance from say 0.99 to 1 takes absolute priority over increasing lower chances. I think that would be too extreme, since it seems to me that improving the chances of those whose survival chances are otherwise very low justifies running some risk of saving fewer people than we could. A too great increase in marginal value when chances approach certainty would also undermine my justification of using a lottery in the example with the two patients, if we assume that giving a patient a kidney would ensure their survival. But how much the marginal moral value of survival chances should increase close to the upper end, and similarly how great risk of saving fewer people than we could we should tolerate, is not something I can say with any precision.

The suggestion that the marginal moral value of survival chances starts increasing as they approach certainty, even though it decreases up to that point, has the added benefit of fitting with how many people seem to reason about how valuable chances are to them.<sup>48</sup> Therefore, if we want the "social planner" to satisfy the *ex ante* Pareto principle, while maximising the moral value of chances, we may have to assume that the marginal moral value of chances increases when the chances get close to certainty.<sup>49</sup> However, the suggestion



has anti-egalitarian implications in some situations. For instance, it implies that in a choice between raising one person's survival chance from 0.59 to 0.6, or instead raising another person's survival chances from 0.99 to 1, we should do the latter (assuming that other things are equal).<sup>50</sup> That may seem morally unappealing. But maybe the reduced risk of saving fewer people than one could justify this anti-egalitarian implication.

## 9 Concluding remarks

In conclusion, I hope to have now established that there is a principled justification of the identified person bias; that this justification satisfies plausible desiderata on any such justification; and that it moreover fits with a natural justification of the use of lotteries to distribute indivisible goods. Now, I should perhaps reiterate that improving people's survival chances is not the only way to increase moral value. My claim is that, *other things being equal* (recall p. 6), moral value is higher the more equally survival chances are distributed. But other things are not always equal. For instance, improving the survival chances of those whose chances are low may sometimes conflict with improving the situation of those who are badly off more generally. In such situations, one may have to trade the moral value of improving the survival chances of those with low chances against other moral values.<sup>51</sup> To make such trade-offs, we need to be able to reason about how changes in survival chances affect moral value. The account defended in this paper can help us with that.

## Notes

1. For discussions of this issue, see for instance [Diamond \(1967\)](#), [Broome \(1984, 1991a,b\)](#), [Stone \(2007, 2011\)](#), [Stefánsson \(2015\)](#), [Stefánsson and Bradley \(2015\)](#), [Henning \(2015\)](#), [Vong \(2015, 2020\)](#) and [Nissan-Rozen \(2017, 2019\)](#).
2. As I further discuss on page 7, my aim is neither to explain nor describe how people actually reason in cases like these. People may in fact think we should prioritise the already trapped miner for reasons that have nothing to do with the identified person bias. My aim is to *justify* certain moral judgements; whether the justification fits with how people actually reach these judgements is an empirical question that I shall not try to answer.
3. For an overview, see [Cohen et al. \(2015\)](#).
4. Examples include [Hare \(2012\)](#), [Daniels \(2015\)](#), [Frick \(2015a,b\)](#), [Slote \(2015\)](#), and [Zameska \(2022\)](#).

- (Hare, 2012 and Daniels, 2015 however do not try to justify saving identified lives at the expense of a *greater* number of statistical lives. I shall get back to this issue in section 4.)
5. For instance, Adler (2015), Brock (2015), Eyal (2015), and Otsuka (2015b).
  6. The only potential exception I know of is Nissan-Rozen (2019). However, he does not use this idea to justify the identified person bias. The idea that chances have decreasing marginal *utility* has however been used to explain observed (and seemingly reasonable) aversion to both risk and ambiguity (the latter being uncertainty about probabilities). See Stefánsson and Bradley (2015, 2019) and also Bradley (2016) (a longer version of which, that includes a discussion of lotteries, is available at <https://personal.lse.ac.uk/bradleyr/pdf/Ellsbergs%20paradox10.pdf>.)
  7. Thanks to Katharina Berndt-Rasmussen for making me see the need to emphasise this.
  8. The same can of course be said about, say, justifying equality increasing policies based on the decreasing marginal utility of money: such a justification is consistent with the fact that money *sometimes* has increasing marginal utility.
  9. Here I am setting aside the fact that it is *possible* that not improving safety at the mine will result in considerably more than one fatality. The value of avoiding catastrophic concentrations of fatalities may, in this case, have to be weighed against the value of improving the survival chances of those whose survival chances are low.
  10. Thanks to Katharina Berndt-Rasmussen for making me see the need to emphasise this.
  11. Both referees for this journal made this observation.
  12. I thank a referee for encouraging me to clarify this.
  13. In other words, I want to allow for a violation of what Stefánsson and Bradley (2015) call *Chance Instrumentalism*.
  14. Thanks to a referee for pressing me on this.
  15. This example was suggested to me by Campbell Brown.
  16. Loewer (2001) develops an account of deterministic chance that is also inspired by Lewis, and according to which the above two conditions could also be satisfied.
  17. Thanks to Nir Eyal for suggesting that I make this explicit.
  18. I am grateful to Barry Loewer for a helpful discussion of Humean chances.
  19. This means that “identified” in the identified person bias that I can justify should not be understood epistemically; that is, the issue is not what some person as a matter of fact has identified or can identify, but rather what is in principle identifiable.
  20. Sridharan (2022) has recently defended a notion of “procedural chances”, which he applies in a novel and very interesting way to the identified person bias. I will not consider his proposal in detail, but two remarks might be worth making. First, the (in his view) intuitively correct judgements that his account delivers are also implied by the decreasing marginal moral value of survival chances. Second, unlike my account, Sridharan’s account cannot help us reason about *how many* expected statistical lives to sacrifice in those cases where the account does imply the identified person bias. I return to this issue in section 4.

21. Here is a quantitative illustration. To say that chances have decreasing marginal moral value is formally equivalent to saying that there is some concave function on chance-values that represents their moral value. The square root is a simple concave function that is often used for illustrative purposes. Now suppose that the 100 not yet trapped miners (in the first mine example) have a survival chance of 0.8 each if we don't improve safety at the mine. The square root function then implies that the aggregate gain in moral value if we improve their survival chances by one percentage point each is 0.55728 while the aggregate gain in moral value if we improve their survival chance by 1.0001 percentage point each is 0.55733; however, the moral value gained if we save the already trapped miner is 1. So, in either case, we should save the already trapped miner (i.e., we get a result that is stronger than Daniels'). However, the square root function implies that the aggregate moral value gained if we improve the survival chances of 10 million miners by ten percentage points each, from 0.8, is 542 561, that is, more than 500 thousand times greater than the moral value that is gained by saving a single trapped miner. So, the same feature of the value function that delivered a result stronger than Daniels' avoids the extreme implication of Frick's justification.
22. For a more detailed discussion of this latter point, that is, that this framework allows for different attitudes to risks and chances depending on what good (or bad) is at stake, see Stefánsson and Bradley (2019: 95-97).
23. By focusing on chances, I may moreover avoid the problems that prioritarians have in accounting for *negative* welfare (as a referee pointed out to me). For a discussion of such problems, see, for instance, Brown (2007).
24. Following the convention in decision theory, I am calling the relevant contingencies "events" rather than "states" when they do not resolve all uncertainty. (See, e.g., Savage, 1954.)
25. For instance, if we again use the square root function to formalise the decreasing marginal moral value of survival chances (recall endnote 21), then the expected moral value of C is roughly 1.4 whereas the expected moral value of D is roughly 1.38.
26. My view may nevertheless be vulnerable to a related criticism: just like *ex ante* prioritarianism, it violates what is called *State-wise dominance*. (I thank a reviewer for pressing me on this.) Recall, for instance, the example with the two patients who both need a hospital's only kidney, with which I opened this article. Given the assumptions of the example, it seems plausible that an outcome where one patient survives as a result of receiving the kidney is just as good as the outcome where the other patient does so. However, I also suggested that holding a lottery to decide which patient gets the kidney is strictly better than giving the kidney to either patient without holding such a lottery. So, on one way of describing these alternatives, the lottery is strictly better than giving the kidney to either patient without holding the lottery, even though the *outcome* is guaranteed to be equally good whichever alternative we choose. And that violates State-wise dominance. Alternatively, one could redescribe the outcomes of the lottery—for instance, by including chances in their descriptions—in a way that dissolves the conflict with State-wise dominance (cf. Broome, 1991b). In Stefánsson (2023) I argue that such a description is appropriate for some, but not all, purposes.
27. The worries raised in the remainder of this section assume that prioritarianism is applicable even when there are no conflicting claims. I will return to this issue in section 7.
28. If Ann's "interest" is that her expected welfare is maximised, then choosing F will in fact be in her interest. This may be so even if it is not in *everyone's* interest that their expected welfare is maximised; for Ann may be *risk neutral* with respect to her own welfare.
29. This follows e.g. from the calculations reported in endnote 25.

30. I am grateful to Alex Voorhoeve for pressing this objection.
31. It is worth noting, however, that some have pointed out that there is a notion of *absolute* fairness too. See, e.g., [Hooker \(2005\)](#).
32. Versions of this question are discussed by, for instance, [Broome \(1984\)](#), [Nissan-Rozen \(2019\)](#), and [Sridharan \(2022\)](#).
33. It is worth noting that if we add that prioritarianism is *only* applicable when people have competing claims, then we can avoid the counterintuitive implications of prioritarianism in single person cases that I discussed in section 5.
34. The idea that lotteries allow a distributor to satisfy—to the greatest extent possible—competing claims to an indivisible good was suggested by [Broome \(1984, 1991a\)](#) and recently made precise by [Nissan-Rozen \(2019\)](#). I should note that the latter offers an explanation for why we should not hold a second lottery in the kidney example that differs somewhat from the explanation I offer here. In particular, [Nissan-Rozen \(2019, 490-492\)](#) demonstrates that his account, which also combines the idea of decreasing marginal value of chances with Broome’s competing claims view, can make sense of the intuitive idea that we should not hold a second lottery, since after the first lottery one patient has ownership of the kidney, which would make the second lottery *unfair*. While the account I am proposing here isn’t rich enough to deliver that explanation, I do think that our explanations are complementary rather than mutually inconsistent. But comparing our accounts in more detail would take me too far from the core aim of this paper, which is to justify the identified person bias.
35. Many thanks to Eyal for a very helpful correspondence about his example.
36. Therefore, I contend, the situation that arises when the Germans bomb a particular house after one turns off the streetlamp (potentially different from the house they would have bombed had the lamp not been turned off), is akin to the situation after the outcome of the coin toss is determined in the two patients example, rather than to the situation that arises once a miner finds himself trapped in the mine. Recall that we are assuming that a house will be bombed whatever we do. In the first two examples, unlike the third, what has therefore changed is that a value maximising process—that is, one that distributes a fixed sum total of survival chances as equally as possible—has taken place. (Thanks to a referee for pressing me on this issue.)
37. Unless, perhaps, the reason that one of other miners is immediately put at risk is that the rescue mission itself puts them at risk. But then other things are not equal, and we are back at the first response: the fact that we should not actively risk people’s lives might pull all things considered moral betterness in one direction while the decreasing marginal moral value of chances pulls it in the opposite direction.
38. It may however be worth noting that due to chances’ upper and lower bounds, my account here won’t face the extreme costs that [Nebel and Stefánsson \(2023\)](#) identify for *ex post* egalitarianism and *ex post* prioritarianism.
39. I thank Anna Mahtani for an informative correspondence about this issue.
40. There might of course be *some* reason in some circumstances to pay to avoid getting the information, for instance, if we think that we will find it impossible to administer the vaccine once we know whom it will kill.
41. It is also possible that the information gained would convince us that we should switch to the first vaccine; for instance, if it turns out that the family of the child who will die if we choose

- the second vaccine has already been harmed, say, by having already lost a child (see, e.g., [Vong and Levinson, 2020](#): 920).
42. I thank a referee for pressing me on this.
  43. Many thanks to Greg Bognar raising this objection.
  44. As those familiar with [Hare \(2012\)](#) will have noticed, the same is true of his three cases.
  45. I am grateful to a referee for pressing me on this issue, and for making me appreciate the importance of responding to it properly.
  46. Giving the four a survival chance of 0.94 while the one only gets a 0.06 chances results in a total moral value of:  $4 \cdot \sqrt{0.94} + \sqrt{0.06} \approx 4.123$ . By contrast, giving all five people a 0.5 chance results in a total moral value of:  $5 \cdot \sqrt{0.5} \approx 3.536$ .
  47. This would in some ways be analogous to Broome's suggestion that in cases like these, *fairness* demands that we give each person an equal chance but the *goodness* of saving a greater number of lives is more important than fairness. For a discussion, see, e.g., [Broome \(1998\)](#).
  48. For discussion see, e.g., [Stefánsson and Bradley \(2015\)](#) and [Stefánsson \(2020\)](#).
  49. To take an example, suppose that the social planner has the choice between spending one million pounds on improving the survival chances of a group of people when their chances are 0.59 or instead spend that money on improving the same people's survival chances at some other time when their chances are 0.99. In light of the above observation, we might predict that each person will prefer the latter, but if the marginal value of survival chances is everywhere decreasing, and the social planner maximises moral value, then they will prefer the former, in violation of *ex ante* Pareto.
  50. In other words, the social planner does not in all cases satisfy a Pigou-Dalton principle for chances. For a discussion of that principle, see [Stefánsson \(2023\)](#).
  51. [Nissan-Rozen \(2019\)](#) makes a similar point.

## References

- Adler, M. (2012). *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. Oxford University Press.
- Adler, M. (2015). Welfarism, equity, and the choice between statistical and identified victims. In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.
- Bradley, R. (2016). Ellsberg's paradox and the value of chances. *Economics and Philosophy*, 32:231–248.
- Brock, D. W. (2015). Identified versus statistical lives: Some introductory issues and arguments. In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified*

*fied versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.

Broome, J. (1984). Uncertainty and fairness. *Economic Journal*, 94:624–32.

Broome, J. (1991a). Fairness. *Proceedings of the Aristotelian Society*, 91:87–101.

Broome, J. (1991b). *Weighing Goods*. Basil Blackwell.

Broome, J. (1998). Kamm on fairness. *Philosophy and Phenomenological Research*, 58:955–961.

Brown, C. (2007). Prioritarianism for variable populations. *Philosophical Studies*, 134:325–361.

Cohen, I. G., Daniels, N., and Eyal, N. (2015). Statistical versus identified persons: An introduction. In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.

Daniels, N. (2015). Can there be moral force to favoring an identified over a statistical life? In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.

Diamond, P. (1967). Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: Comment. *Journal of Political Economy*, 75:765–766.

Eyal, N. (2015). Concentrated risk, the Coventry Blitz, Chamberlain's cancer. In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.

Fleurbaey, M. and Voorhoeve, A. (2013). Decide as you would with full information! An argument against ex ante Pareto. In Norheim, O., Hurst, S., Eyal, N., and Wikler, D., editors, *Inequalities in Health: Concepts, Measures, and Ethics*. Oxford University Press.

Frick, J. (2015a). Contractualism and social risk. *Philosophy and Public Affairs*, 43:175–223.

Frick, J. (2015b). Treatment versus prevention in the fight against HIV/AIDS and the problem of identified versus statistical lives. In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.

- Frigg, R. and Hoefer, C. (2010). Determinism and chance from a Humean perspective. In Dieks, D., Gonzalez, W., Hartmann, S., Weber, M., Stadler, F., and Uebel, T., editors, *The Present Situation in the Philosophy of Science*, pages 351–271. Springer.
- Frigg, R. and Hoefer, C. (2015). The best Humean system for statistical mechanics. *Erkenntnis*, 80:551–574.
- Hare, C. (2012). Obligations to merely statistical people. *Journal of Philosophy*, 109:378–390.
- Henning, T. (2015). From choice to chance? Saving people, fairness, and lotteries. *Philosophical Review*, 124:169–206.
- Hoefer, C. (2007). The third way on objective probability: A sceptic’s guide to objective chance. *Mind*, 116:449–496.
- Hooker, B. (2005). Fairness. *Ethical Theory and Moral Practice*, 8:329–352.
- Lewis, D. (1980). A subjectivist’s guide to objective chance. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability*. University of California Press.
- Loewer, B. (2001). Determinism and chance. *Studies in History and Philosophy of Science Part B*, 32:609–620.
- Nebel, J. M. and Stefánsson, H. O. (2023). Calibration dilemmas in the ethics of distribution. *Economics and Philosophy*, 39:67–98.
- Nissan-Rozen, I. (2017). How to be an ex-post egalitarian and an ex-ante Paretian. *Analysis*, 77:550–558.
- Nissan-Rozen, I. (2019). The value of chance and the satisfaction of claims. *Journal of Philosophy*, 116:469–493.
- Otsuka, M. (2015a). Prioritarianism and the measure of utility. *Journal of Political Philosophy*, 23:1–22.
- Otsuka, M. (2015b). Risking life and limb: How to discount harms by their improbability. In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.
- Otsuka, M. and Voorhoeve, A. (2009). Why it matters that some are worse off than others: An argument against the priority view. *Philosophy and Public Affairs*, 37:171–199.
- Otsuka, M. and Voorhoeve, A. (2018). Equality versus priority. In Olsaretti, S., editor, *Oxford Handbook of Distributive Justice*. Oxford University Press.



- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Parfit, D. (2002/1991). Equality or priority? In Clayton, M. and Williams, A., editors, *The Ideal of Equality*. Palgrave.
- Rabinowicz, W. (2002). Prioritarianism for prospects. *Utilitas*, 14:2–21.
- Savage, L. (1972/1954). *The Foundations of Statistics*. Dover Publication.
- Schaffer, J. (2007). Deterministic chance? *British Journal for the Philosophy of Science*, 58:113–140.
- Schelling, T. C. (1968). The life you save may be your own. In Chase, S. B., editor, *Problems in Public Expenditure Analysis*, pages 127–162. The Brookings Institution.
- Slote, M. (2015). Why not empathy? In Cohen, I. G., Daniels, N., and Eyal, N., editors, *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Oxford University Press.
- Sridharan, V. (2022). Procedural chances and the equality of claims. *Philosophy and Phenomenological Research*.
- Stefánsson, H. O. (2015). Fair chance and modal consequentialism. *Economics and Philosophy*, 31:371–395.
- Stefánsson, H. O. (2020). Gambling with death. *Topoi*, 39:271–281.
- Stefánsson, H. O. (2023). In defence of Pigou-Dalton for chances. Unpublished manuscript.
- Stefánsson, H. O. and Bradley, R. (2015). How valuable are chances? *Philosophy of Science*, 82:602–625.
- Stefánsson, H. O. and Bradley, R. (2019). What is risk aversion? *British Journal for the Philosophy of Science*, 70:77–102.
- Stone, P. (2007). Why lotteries are just. *Journal of Political Philosophy*, 15:276–295.
- Stone, P. (2011). *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press.
- Temkin, L. S. (1993). *Inequality*. Oxford University Press.
- Vong, G. (2015). Fairness, benefiting by lottery and the chancy satisfaction of moral claims. *Utilitas*, 27:470–486.
- Vong, G. (2020). Weighing up weighted lotteries: Scarcity, overlap cases, and fair inequalities of chance. *Ethics*, 130:320–348.

- Vong, G. and Levinson, M. (2020). The ethics of risk displacement in research and public policy. *Bioethics*, 34:918–922.
- Voorhoeve, A. and Fleurbaey, M. (2012). Egalitarianism and the separateness of persons. *Utilitas*, 24:381–398.
- Zameska, J. A. (2022). An uncertainty argument for the identified victim bias. *Journal of Applied Philosophy*, 39:504–518.