

Priority or Equality for Possible People?*

Alex Voorhoeve and Marc Fleurbaey

Suppose that you must make choices that may influence the well-being and the identities of the people who will exist, though not the number of people who will exist. How ought you to choose? This article answers this question. It argues that the currency of distributive ethics in such cases is a combination of an individual's final well-being and her expected well-being conditional on her existence. It also argues that this currency should be distributed in an egalitarian, rather than a prioritarian, manner.

I

Suppose that you are a morally motivated stranger who must make choices that may influence the well-being and the identities of the people who will exist, though not the number of people who will exist. How ought you to choose? This article addresses this question.

By way of introduction, consider the following case.¹

Interpersonal, One Future Person Case: Either Ann or Bob will come into existence. Both are equally likely to exist. Their chances of existing are independent of your actions. Unless you intervene, the well-being of whoever exists will be low (a well-being of 10). At negligible cost to yourself, you can improve the life of whoever exists by choosing one of the prospects de-

* This article was presented at Carnegie Mellon, Erasmus University Rotterdam, the National Institutes of Health, the University of Oxford, Princeton University, Stockholm University, the University of Amsterdam, the University of Warwick, and to the Society of Applied Philosophy. We thank those audiences and Luc Bovens, Richard Bradley, Susanne Burri, Benjamin Ferguson, Benjamin Lange, Silvia Milano, Joseph Millum, Fabienne Peter, Wlodek Rabinowicz, Thomas Rowe, Richard Rowland, Hlynur Orri Stefansson, Shlomi Segall, Patrick Tomlin, Peter Vallentyne, the editors of this journal, and two referees for comments. We are especially indebted to Michael Otsuka for extensive comments and discussion. Alex Voorhoeve is grateful for a Faculty Fellowship of the Princeton University Center for Human Values and for support from the British Arts and Humanities Research Council through grant AH/J006033/1.

1. This case is adapted from Michael Otsuka, "Prioritarianism and the Separateness of Persons," *Utilitas* 24 (2012): 365–80, 369.

Ethics 126 (July 2016): 929–954

© 2016 by The University of Chicago. All rights reserved. 0014-1704/2016/12604-0003\$10.00

scribed in table 1. (*S1* and *S2* are states of the world; the numbers in parentheses are these states' probabilities, rationally calculated given all the available evidence. Nonexistence is represented by an empty column. Since the cost to you is negligible and the good you can do is great, we will assume throughout that the "no-intervention" option is obviously ineligible. We therefore represent only the two possible ways of doing good.)

This case represents, in abstract form, scenarios in which one must make choices to improve the well-being of future generations without knowing who will be born. One such scenario is this. A couple is living in a place that has poor prospects for children. This couple will have one child, which is equally likely to be a girl or a boy. You can help them relocate to a better place. You must either (prospect A) help them move to a place where the future child will have a well-being of 60 (a moderately good life) independently of whether it is a girl or a boy, or instead, (prospect B), help them move to a place where, if the child is a girl, she will have a well-being of 40 (an okay life), and if the child is a boy, he will have a well-being of 80 (a very good life).

Another scenario is this. You must now make provision to help a future generation adapt to global warming. It is uncertain how much the Earth will warm, and since the degree of warming will determine who will be born, the way this uncertainty resolves will determine who is born. Suppose that warming will be either three or two degrees Celsius, that both are equally probable, and that these probabilities are independent of your actions. If it is three degrees, all members of the future generation will have the quality of life of Ann; if it is two degrees, they will all have the quality of life of Bob (there is, let us suppose, no intragenerational inequality). You must either (prospect A) do what will provide a large benefit to the future generation no matter whether the Earth warms by three or two degrees, or instead (prospect B) do what will provide only a moderate benefit if the Earth warms by three degrees, but a very large benefit to the future generation if the Earth warms by two degrees.

TABLE 1
FINAL WELL-BEING FOR THE INTERPERSONAL,
ONE FUTURE PERSON CASE

	<i>S1</i> (0.5)		<i>S2</i> (0.5)	
	Ann	Bob	Ann	Bob
Prospect A	60			60
Prospect B	40			80

Considering Ann's and Bob's potential fates in isolation from how anyone else fares, how should you choose? To us, it is clear that you ought to choose prospect A. It is more important to provide Ann with the additional benefit of existing at well-being level 60 rather than at 40 if *S1* is the case than to provide Bob with the additional benefit of existing at 80 rather than 60 if *S2* is the case, because the former would improve Ann's well-being from a lower level. In this article, we attempt to find a view that justifies this verdict, that yields acceptable answers in other cases, and that has a plausible rationale in respect for both the unity of the individual and the separateness of persons.

In order to do so, we tackle two central issues in the ethics of distribution, which are relevant in both familiar risky, fixed-identity cases and in less extensively studied risky, same-number, nonidentity cases. The first issue is the currency of well-being-based distributive ethics. Suppose (as we do in this article) that one should be concerned with a person's well-being. This still leaves open the following questions. Should one be concerned with only her final well-being, with only her expected well-being, or with some combination of the two? And, if expected well-being is of concern, how should one assess the expected well-being of a person (such as Ann in the Interpersonal, One Future Person Case) who merely has some chance of existing? In response to these questions, we will argue that in risky, fixed-identity cases, the currency of distributive ethics is a hybrid of final and expected well-being. We will also argue that in risky, same-number, nonidentity cases, when considering a merely possible person, one should be concerned with a hybrid of her final well-being and her expected well-being conditional on her existence.

The second issue is how one ought to distribute this currency. The judgment that one ought to choose A in the Interpersonal, One Future Person Case challenges two familiar answers to this question. This judgment is inconsistent with a form of utilitarianism which mandates maximizing expected total well-being, because the expected total well-being of A is identical to that of B, so that, on this utilitarian view, B is also a permissible choice. Moreover, the judgment that one should choose A is not explained by a common egalitarian view.² On this view, it is bad that some individuals who have existed, currently exist, or will exist are worse off than others (through no choice or fault of theirs).³ But in this case, no such inequality obtains, since of the possible individuals whose fate we are considering, only one will ever exist. Egalitarians who wish to

2. *Ibid.*, 370.

3. G. A. Cohen, "On the Currency of Egalitarian Justice," *Ethics* 99 (1989): 906–44; Larry Temkin, *Inequality* (Oxford: Oxford University Press, 1993); and Richard Arneson, "Postscript to 'Equality and Equal Opportunity for Welfare,'" in *Equality: Selected Readings*, ed. L. Pojman and R. Westmoreland (Oxford: Oxford University Press, 1997), 238–41.

justify a requirement to choose A therefore need to appeal to something other than the badness of such inequality.

By contrast, prioritarianism, according to which an increment in an individual's well-being matters more the lower the level of well-being from which this increment takes place, readily arrives at the right answer. On this view, you should choose A because the potential increase in Ann's well-being from 40 to 60 has greater expected prioritarian value than the potential increase in Bob's well-being from 60 to 80. This case therefore seems to provide a reason to regard prioritarianism as at least part of the correct theory of distributive ethics.⁴

In the light of its apparent success in this case, in this article, we develop the most plausible version of prioritarianism for risky, nonidentity, same-number cases. We will argue, however, that there is no reason to endorse this form of prioritarianism as either part or the whole of one's theory of distributive ethics. For, we argue, there is an even more plausible pluralist egalitarian view in which this form of prioritarianism plays no part. This egalitarian view arrives at the right conclusion in this and other cases and has two merits that prioritarianism lacks. First, it fully respects the unity of the person. Second, it is more willing to accept a costly transfer (one that involves some loss in expected total well-being) to a person with worse prospects when this reduces inequality between people who both exist than when it does not reduce such inequality.

Before proceeding, a comment on some of our assumptions will be helpful. Throughout, we assume that orthodox decision theory applies, according to which under risk, one ought to maximize the expectation of the value with which one is concerned (so that a utilitarian should maximize expected total final well-being, a final-well-being prioritarian should maximize the expectation of total priority-weighted final well-being, etc.). We also assume a cardinal, interpersonally comparable measure of well-being, or prudential value, derived from idealized preferences satisfying the Von Neumann–Morgenstern axioms. On this measure, a prospect has higher expected well-being for a person just in case it would be preferred for that person's sake after rational and calm deliberation with all pertinent information while attending to his self-interest only. One prospect has the same expected well-being as another for a person just in case such deliberation would yield indifference between the two prospects.⁵

4. For example, it appears to give pluralist egalitarians a reason to endorse prioritarianism in addition to egalitarianism. For such mixed "egalitarian-prioritarian views," see Martin Peterson and Sven Ove Hansson, "Equality and Priority," *Utilitas* 17 (2005): 299–309; and Shlomi Segall, "In Defense of Priority (and Equality)," *Philosophy, Politics and Economics* 14 (2015): 343–64.

5. This measure does not presuppose any particular view on what well-being is. One might maintain that well-being consists of something other than preference satisfaction

Furthermore, we stipulate that well-being level 0 is a quality of life such that, from the perspective of the idealized preferences of a person living that life, it is a matter of indifference that he lives that life or never existed.⁶ A life with positive well-being is one that the person who lives it would rationally prefer to never existing; a life with negative well-being is one that he would rationally disprefer to never existing.

We must also clarify what we mean by the term “currency of distributive ethics,” or object of distributive concern. We will consider various specifications of this currency: final well-being, expected well-being, and a mix of the two. We stipulate that across all specifications, the concept exhibits the same tight connection we have just assumed between well-being and prudential value. Across all specifications, it therefore tracks prudential value in the following way: when a prospect has greater expected prudential value for an individual, then it gives him a higher expected amount of the currency of distributive ethics; when it has equal expected prudential value for an individual, then it gives him an equal expected amount of this currency; and when it has lower expected prudential value, then it gives him a lower expected amount of this currency.

Finally, we assume that the distributor’s actions will never lower anyone’s well-being compared to a baseline of inaction (without the distributor’s intervention, the well-being of whoever exists will be low; the distributor’s choices may improve their prospects). We also stipulate that all lives that might come about have positive well-being. We do so in order to focus solely on the distribution of chances of benefits and goods, rather than having to balance them against risks of harms or evils, which a decision maker may have special reason to avoid.

This article is organized as follows. Section II draws on recent work in distributive ethics to develop the most plausible version of prioritarianism for risky, fixed-identity cases. Section III does so for egalitarianism. These sections conclude that the most plausible version of each takes the object of distributive concern to be a hybrid of final and expected well-being. Sections IV and V consider different ways of extending these hybrid views to risky, fixed-number, nonidentity cases. They conclude that in such cases, the currency of distributive ethics is a hybrid of an individual’s final well-being and her expected well-being conditional on her existence. Section VI argues that egalitarianism for this currency is superior to prioritarianism. Section VII concludes.

and hold that the specified idealized preferences fully track the magnitude of this other thing. See Michael Otsuka and Alex Voorhoeve, “Why It Matters that Some Are Worse Off than Others,” *Philosophy and Public Affairs* 37 (2009): 171–99, 172–73 n. 3.

6. In this, we follow Matthew W. Adler, *Well-Being and Fair Distribution* (Oxford: Oxford University Press, 2012), 219–20.

II

In conditions of certainty, the prioritarianism we consider is based on the following three ideas.

Diminishing Marginal Value: Increments in a person's well-being have positive but diminishing marginal moral value—an increment that takes place from a lower level receives a higher “priority weight” than an increment that takes place from a higher level;

Separability: The moral value of an increment in a person's well-being depends only on his level of well-being and not on how anyone else fares; and

Maximization: We ought to maximize the sum-total of moral value, which is the sum of priority-weighted well-being.⁷

How should we extend this view to risky contexts? One approach is “final well-being prioritarianism.” This takes the currency of distributive ethics to be identical to final well-being. It therefore applies priority weights to each person's holdings of this currency. Under risk, it directs one to maximize the expected sum of priority-weighted final well-being.⁸ The following Intra- versus Interpersonal Case shows that this approach is problematic, however.⁹ This case contrasts two scenarios.

Intrapersonal Scenario: In this scenario, you must choose between A and B (Intrapersonal) outlined in table 2 ($d > 0$).

Interpersonal Scenario: In this scenario, you must choose between A and B (Interpersonal) outlined in table 2 ($d > 0$).

In the Intrapersonal Scenario, only Chelsey's well-being is at stake in the choice between A and B. The latter exposes Chelsey to a risk of ending up less well off than she might, but also gives her a chance of ending up much better off than she otherwise would. Since d is positive, B maximizes Chelsey's expected well-being. Given our assumed measure

7. See Derek Parfit, “Equality or Priority?” in *The Ideal of Equality*, ed. M. Clayton and A. Williams (Basingstoke: Palgrave Macmillan, 2002), 81–125, on 104–5; and Adler, *Well-Being and Fair Distribution*. We do not discuss an alternative form of prioritarianism proposed by Andrew Williams, “The Priority View Bites the Dust?” *Utilitas* 24 (2012): 315–31, which violates Separability.

8. This form of prioritarianism is proposed in Wlodek Rabinowicz, “Prioritarianism for Prospects,” *Utilitas* 14 (2002): 2–21; and Adler, *Well-Being and Fair Distribution*.

9. See Otsuka and Voorhoeve, “Why It Matters”; and Alex Voorhoeve and Marc Fleurbaey, “Egalitarianism and the Separateness of Persons,” *Utilitas* 24 (2012): 381–98.

TABLE 2
FINAL WELL-BEING FOR THE INTRA- VERSUS INTERPERSONAL CASE

	<i>S1</i> (0.5)		<i>S2</i> (0.5)	
	Chelsey	Dave	Chelsey	Dave
Prospect A	60	60	60	60
Prospect B (Intrapersonal)	40	60	$80 + d$	60
Prospect B (Interpersonal)	40	60	60	$80 + d$

of well-being, someone solely concerned with Chelsey’s interests would therefore prefer B. This means that no matter how things turn out, if you were to choose B, you would be able to offer her the following prudential justification: “I did the best I could for you, given the information I had at the time.” This provides you with a strong reason to choose B.¹⁰

In contrast, in the Interpersonal Scenario, there is a conflict of interest between Chelsey and Dave. Choosing B would give Dave a chance at a gain at Chelsey’s expense, while making her worse off than Dave for sure.

These differences between these scenarios matter. For any given positive d , it is easier to justify exposing Chelsey to a 50 percent chance of doing less well than she might for the sake of giving her a 50 percent chance at a great benefit than it is to justify exposing Chelsey to a 50 percent chance of doing less well than she might for the sake of giving Dave a 50 percent chance at a great benefit, thereby ensuring that he will be better off than her no matter what happens. There will therefore be some d for which B is permissible in the Intrapersonal Scenario but for which B is not permissible in the Interpersonal Scenario.

Final well-being prioritarianism cannot account for this difference in justifiability. The anonymized distribution of final well-being under B is the same in either scenario—there is either (in *S1*), one person at 40 and one at 60, or (in *S2*), one person at 60 and one at $80 + d$. It follows that, on this view, for every d for which B is at least as good as A in the Intrapersonal Scenario, B is also at least as good as A in the Interpersonal Scenario. In failing to accommodate the difference in justifiability between the risky options in these scenarios, final well-being prioritarianism fails to respect the difference between the unity of the individual

10. We do not assume that this is a decisive reason. Indeed, in our view, the inequality present in this case means that for some sufficiently small d , one ought to choose A. See Sec. III for discussion.

(which gives us reason to make purely intrapersonal trade-offs for a person's sake) and the separateness of persons (which requires that in interpersonal trade-offs, we give greater weight to the interests of the worse off).

A different version of prioritarianism avoids this objection. "Expected well-being prioritarianism" takes the currency of distributive ethics to be identical to expected well-being. It therefore applies priority weights to each person's "holdings" of the currency of expected well-being. It directs one to maximize the sum of priority-weighted expected well-being. In intrapersonal trade-offs, it requires maximizing expected well-being; in interpersonal trade-offs, it gives priority to whoever has lower expected well-being. It therefore requires choosing B for every positive d in the Intrapersonal Scenario, but also holds that for some sufficiently small, positive d , you ought not to choose B in the Interpersonal Scenario. It therefore recognizes the difference in the justifiability of these prospects.

However, the Inversely Correlated Case described in table 3 illustrates that expected well-being does not exhaust the currency of prioritarian ethics.

Both prospects give Edie and Fred the same expected well-being. Expected well-being prioritarianism is therefore indifferent between them. By contrast, it is clear to us that you ought to make the life of whoever ends up worst off as good as possible in this case, in which this can be accomplished without loss in total well-being. One way of arriving at this conclusion is to note that while there are no conflicts of interest in expected well-being in this case, there are conflicts of interest in terms of final well-being. In *S1*, A is in Edie's final well-being interest, while B is in Fred's. In *S2*, A is in Fred's final well-being interest, while B is in Edie's. In this interpersonal trade-off in final well-being, we submit that the separateness of persons gives you reason to assign extra importance to improving the fate of the less well-off person.¹¹

Together, the Intra- versus Interpersonal Case (table 2) and the Inversely Correlated Case (table 3) suggest that prioritarians should take the currency of distributive ethics to be a combination of expected well-being and final well-being. A person's final well-being matters simply because this is how well her life goes. A person's expected well-being matters because it determines how well an alternative advances her interests assessed at the moment of decision (when the decision maker has only probabilistic knowledge of the true state of the world). Even if (as we

11. For further discussion, see Adler, *Well-Being and Fair Distribution*, 495–523; Marc Fleurbaey and Alex Voorhoeve, "Decide as You Would with Full Information! An Argument against Ex Ante Pareto," in *Inequalities in Health: Concepts, Measures, and Ethics*, ed. N. Eyal, S. Hurst, O. Norheim, and D. Wikler (Oxford: Oxford University Press, 2013), 113–28; and Johann Frick, "Uncertainty and Justifiability to Each Person: A Response to Fleurbaey and Voorhoeve," in the same volume, 129–46.

TABLE 3
FINAL WELL-BEING IN THE INVERSELY CORRELATED CASE

	S1 (0.5)		S2 (0.5)	
	Edie	Fred	Edie	Fred
Prospect A	60	60	60	60
Prospect B	40	80	80	40

assume) a chance at a benefit that does not come good does not improve a person’s final well-being, it remains the case even after the chance has come to nought that the provision of this chance was something to be sought for her sake at the time of decision. As the Intra- versus Interpersonal Case (table 2) demonstrates, whether an alternative could be chosen for a person’s sake is an important consideration. In order to take account of the presence or absence of a prudential justification to a person for the choice of a particular alternative, we must take account of a person’s expected well-being.¹²

How important are the respective contributions of expected and final well-being to the currency of distributive ethics? To answer this question, consider the Expected versus Final Well-being Case outlined in table 4, with $20 > d > 0$.

For $20 > d > 0$, this is a choice between improving the distribution of expected well-being by choosing A (under which each person’s expected well-being is 60) and improving the distribution of final well-being by choosing B (under which Gina’s expected well-being is $80 - d$ and Herb’s is $40 + d$). For a very small, positive d , B is much worse as regards the distribution of expected well-being and only somewhat better as regards the distribution of final well-being. For a very small, positive d , we submit that one therefore ought to prefer A. As d increases to 20, the situation of the worst off in terms of both expected and final well-being in B improves. Moreover, this improvement takes place without loss in total well-being, so that it represents an unambiguous improvement from a prioritarian perspective.¹³ As d increases, there will therefore be a point at which one should be indifferent between A and B. This point indicates the relative importance of expected and final well-being. If one should be indifferent

12. On this point, see also Keith Hyams, “Hypothetical Choice, Egalitarianism, and the Separateness of Persons,” *Utilitas* 27 (2015): 217–39. As outlined in Sec. III, egalitarians have a further reason to favor concern with a person’s expected well-being, namely, that it is one of the determinants of the fairness of a distribution of well-being.

13. The same is true, of course, from an egalitarian perspective. This example can therefore also serve to calibrate the form of egalitarianism proposed in Sec. III.

TABLE 4

FINAL WELL-BEING FOR THE EXPECTED VERSUS FINAL WELL-BEING CASE

	S1 (0.5)		S2 (0.5)	
	Gina	Herb	Gina	Herb
Prospect A	80	40	40	80
Prospect B	$80 - d$	$40 + d$	$80 - d$	$40 + d$

for a small d , then this implies that one should give relatively little weight to expected well-being, since one ought to accept a large worsening in the distribution of expected well-being for the sake of a small improvement in the distribution of final well-being. By contrast, if one should be indifferent only for a large d (close to 20), then this implies that one should give large weight to expected well-being.

Intuitively, it strikes us that one ought to be indifferent for a modest d —one should give more weight to how well people’s lives truly go than to the quality of their prospects. The resulting form of hybrid prioritarianism, which cares for both expected and final well-being while giving greater weight to the latter, is the most plausible prioritarian view because it recognizes a difference between intra- and interpersonal trade-offs.

Nonetheless, this form of prioritarianism has a drawback: in cases in which only one person’s well-being is at issue and inequality is not an issue, it sometimes mandates the choice of an alternative with lower expected well-being for this person.¹⁴ By way of illustration, consider the Intrapersonal, One Future Person Case outlined in table 5, where $d > 0$.

Prospect B uniquely maximizes Io’s expected well-being. Given our assumed measure of well-being, someone solely concerned with Io’s interests would therefore prefer it. This gives you a strong reason to choose this prospect. Moreover, when you consider her prospects in isolation, you have, in our view, no countervailing reason to choose A. We conclude that it is morally better to choose B.

However, for a sufficiently small, positive d , hybrid prioritarianism can mandate A. To see why, consider first a case in which $d = 0$ (contrary to our assumption above). Hybrid prioritarianism then regards A as

14. Michael Otsuka, “Prioritarianism and the Measure of Utility,” *Journal of Political Philosophy* 23 (2015): 1–22, notes this consequence of hybrid prioritarianism. That final well-being prioritarianism has this implication and that this counts against it was first argued by Dennis McKerlie, “Egalitarianism,” *Dialogue* 23 (1984): 223–37, 235. For concurring judgments, see David McCarthy, “Utilitarianism and Prioritarianism II,” *Economics and Philosophy* 24 (2008): 1–33; Otsuka and Voorhoeve, “Why It Matters”; and Williams, “The Priority View Bites the Dust?”

TABLE 5
FINAL WELL-BEING FOR THE INTRAPERSONAL, ONE PERSON CASE

	<i>S1</i> (0.5)	<i>S2</i> (0.5)
	Io	Io
Prospect A	60	60
Prospect B	40	80 + <i>d</i>

superior to B, because while they are equal in terms of priority-weighted expected well-being, A is more valuable in terms of expected priority-weighted final well-being. Indeed, if one gives substantial prioritarian weight to improvements at lower levels of final well-being, A will be considerably better than B. Now let *d* be marginally larger than zero. This gives B a minuscule advantage in terms of expected well-being, but it remains considerably worse than A from the perspective of priority-weighted final well-being. This way in which B is worse therefore outweighs the advantage that B has in terms of priority-weighted expected well-being. Hybrid prioritarianism therefore mandates the choice of A, contrary to Io’s interests.

We regard this as a reason to reject hybrid prioritarianism. This judgment has been much debated. Since we aim to cover new ground, we shall not review this debate in its entirety. Nonetheless, it may be useful to respond to the following common objection to our judgment.¹⁵

In choosing to benefit a person, one should be risk averse¹⁶ in well-being: when comparing a riskless prospect such as A to a risky prospect such as B with a somewhat higher expected well-being for the person but with a chance of ending up less well off than under A, one should favor the riskless prospect. In giving greater weight to increments in final well-being that take place from a lower level, hybrid prioritarianism displays such proper risk aversion.¹⁷

15. For other objections and replies, see Otsuka and Voorhoeve, “Why It Matters,” 185–95; Roger Crisp, “In Defence of the Priority View: A Response to Otsuka and Voorhoeve,” *Utilitas* 23 (2011): 105–8; and Michael Otsuka and Alex Voorhoeve, “Reply to Crisp,” *Utilitas* 23 (2011): 109–14.

16. We are referring here to the Arrow-Pratt conception of risk aversion. See Kenneth J. Arrow, *Aspects of the Theory of Risk-Bearing* (Helsinki: Yrjö Jahnssoonin Säätiö, 1965), lecture 2; and John Pratt, “Risk Aversion in the Small and in the Large,” *Econometrica* 32 (1964): 122–36.

17. See Martin O’Neill, “Priority, Preference, and Value,” *Utilitas* 24 (2012): 332–48; Parfit, “Another Defense of the Priority View,” *Utilitas* 24 (2012): 399–440, 423; Thomas Porter, “In Defence of the Priority View,” *Utilitas* 24 (2012): 349–64; and Luc Bovens, “Concerns for the Poorly Off in Ordering Risky Prospects,” *Economics and Philosophy* 31 (2015): 397–429.

In reply: the assumption that one ought to be risk averse in well-being when deciding on how to benefit another is unsupported. It may acquire some intuitive plausibility from the fact that risk aversion for goods that contribute to well-being is both common and eminently reasonable because these goods typically have diminishing marginal prudential value.¹⁸ For example, suppose you have two ways of benefiting someone for whom money has diminishing marginal prudential value: you can either provide him with \$100,000 for sure or instead expose him to a fifty-fifty gamble between gaining nothing and \$200,000. In this case, when choosing for his sake, you should be risk averse in money and give him \$100,000. You would thereby choose in line with the risk attitude toward money that he would have if he were ideally rational. But unlike ordinary goods, by construction, well-being has constant marginal prudential value. One therefore does not have the reason to be risk averse in well-being that one has to be risk averse in ordinary goods. On the measure of well-being here assumed, a person would rationally prefer that which uniquely maximizes his expected well-being. Why should a third party deciding on this person's fate adopt a different risk attitude than this person would, if ideally rational, adopt on his own behalf?¹⁹

We have argued that despite its strengths, hybrid prioritarianism fails in the Intrapersonal, One Person Case (table 5) to respect the fact that the two potential futures are both Io's, so that one can legitimately approach this choice with an eye to maximizing expected prudential value for her. In the next section, we develop an egalitarian view that does not have this demerit.

III

The preceding section established that the currency of prioritarian ethics is a combination of expected and final well-being. There are good reasons to take this to be the currency of egalitarianism as well. The basic

18. For the hypothesis that some people are risk averse in well-being because they misapply a heuristic which makes perfect sense for everyday goods, see Joshua Greene and Jonathan Baron, "Intuitions about Declining Marginal Utility," *Journal of Behavioral Decision Making* 14 (2001): 243–55.

19. See Otsuka, "Prioritarianism and the Measure of Utility" and David McCarthy, "The Priority View," *Economics and Philosophy* (forthcoming), for an extended defense of the idea that, in a risky, one-person case, the judgment of what is better for an individual's sake should fully determine what is morally better. Otsuka and McCarthy offer arguments that do not rely on the Von Neumann–Morgenstern measure of well-being. If correct, their arguments therefore establish that one should reject prioritarianism even if one assumes a different measure of well-being than we do here. For discussion of the relation between the Von Neumann–Morgenstern measure and prioritarianism, see Hilary Greaves, "Anti-prioritarianism," *Utilitas* 27 (2015): 1–42.

egalitarian idea is that it is bad because unfair that some are worse off than others in terms of final well-being. The unfairness of such inequality is partly determined by individuals' chances of ending up worse off than others. To illustrate, consider a case in which one has to distribute a single, indivisible good between two people, each of whom would fare equally poorly without it and equally well with it. The resulting inequality in final well-being between these people, while unfair, will be less unfair when each person has an equal chance to end up better off than when the one who ends up worse off never has such a chance. In receiving this equal chance, what the first person receives is, in terms of her interests as assessed with the knowledge of the distributor, just as valuable as what the second receives. Moreover, from the perspective of someone solely concerned with the interests of any one of these persons (e.g., a guardian) and who has the same knowledge as this distributor, this chance is equivalent to getting half the well-being value of the disputed good for sure. Because, in these senses, equal chances advance the interests of each person equally, they mitigate (without eliminating)²⁰ the unfairness of unequal final well-being in such a case.²¹

Sensible egalitarians are pluralists—they care about equality and about improving people's well-being. Here, we will focus on one form of pluralist egalitarianism with especially attractive properties known as the "equally-distributed equivalent," or EDE, view.²² To illustrate this view,

20. We therefore disagree with the proposal in Peter Vallentyne, "Brute Luck, Option Luck, and Equality of Initial Opportunities," *Ethics* 112 (2002): 529–57, that one should take expected well-being at birth to be the whole of the currency of egalitarian distributive ethics. This proposal has, in our view, implausible implications in the Inversely Correlated Case of table 3 (in which it is indifferent between A and B) and the Expected versus Final Well-Being Case of table 4 (in which it favors A for every d in the specified range).

21. See Peter Diamond, "Cardinal Welfare, Interpersonal Ethics, and Interpersonal Comparisons of Utility: Comment," *Journal of Political Economy* 75 (1967): 765–66; Arneson, "Postscript"; John Broome, "Fairness," *Proceedings of the Aristotelian Society* 91 (1990): 87–101; and Michael Otsuka, "The Fairness of Equal Chances" (unpublished manuscript, London School of Economics). Note that we do not share Diamond's judgment that one ought to reject the Sure-Thing Principle. Instead, we regard a person's expected well-being as part of the morally relevant features of an outcome. See John Broome, *Weighing Goods* (Oxford: Oxford University Press, 1991), chap. 5; and H. Orri Stefansson, "Fair Chance and Modal Consequentialism," *Economics and Philosophy* 31 (2015): 371–95.

22. The EDE was introduced by Serge-Christophe Kolm, "The Optimal Production of Social Justice," in *Public Economics*, ed. J. Margolis and H. Guitton (London: Macmillan, 1969), 145–200; and Anthony Atkinson, "On the Measurement of Inequality," *Journal of Economic Theory* 2 (1970): 244–63. Under certainty, the social welfare function they proposed is consistent with prioritarianism, since it respects diminishing marginal value, separability, and maximization. The distinctively egalitarian credentials of an EDE social welfare function become apparent when one considers risky cases. See Marc Fleurbaey, "Equality or Priority: How Relevant Is the Distinction?" *Economics and Philosophy* 31 (2015): 203–17, and "Assessing Risky Social Situations," *Journal of Political Economy* 118 (2010): 649–80.

consider a risk-free prospect that gives Arnaldur 40 and Bea 80 for sure. Suppose an egalitarian should be indifferent between this prospect and an alternative risk-free prospect in which Arnaldur and Bea each have 50. Then the latter is the equally distributed equivalent of the former, unequal prospect. When a risk-free prospect generates inequality, its EDE will be less than the average attainment in that unequal prospect. When such a prospect contains no inequality, its EDE is simply the uniform attainment in that prospect.

In risky, fixed-identity cases, on the EDE view, we evaluate each prospect as follows. First, we establish the EDE for each of the possible outcomes (distributions of the hybrid currency) that may result from the choice of the prospect. Then we take the probability-weighted sum of these values. This view then tells us to choose the prospect with the highest expected EDE.

By way of illustration, consider first a case involving a choice between an equal and an unequal prospect, the Inversely Correlated Case (table 3). Prospect A yields the following outcome for sure: Edie has a final well-being of 60 and had (at the moment of decision) an expected well-being of 60; the same is the case for Fred. Since there is no inequality, the EDE of this prospect is simply this: a final well-being of 60 and an expected well-being of 60.

Prospect B has two possible outcomes. In *S1*, it yields a distribution in which Edie has a final well-being of 40 and had an expected well-being of 60 and Fred has a final well-being of 80 and had an expected well-being of 60. In *S1*, it therefore generates inequality in one component of the hybrid currency (final well-being). The EDE of this outcome will therefore be a final well-being of less than 60 and an expected well-being of 60. In *S1*, the outcome of B is therefore less valuable than the outcome of A. Analogous reasoning establishes that in *S2*, the outcome of B is again less valuable than the outcome of A. The egalitarian prospect A therefore yields higher expected value than the inegalitarian B.

Next consider the Intrapersonal, One Person Case (table 5). The EDE of A is, of course, a final well-being of 60 and an expected well-being of 60. For B, we must again consider two possible outcomes. In *S1*, Io has a final well-being of 40 and an expected well-being of $60 + d/2$ (recall that in this case, $d > 0$). Since we are considering Io in isolation, there is no inequality. This therefore is just the EDE of this outcome. In *S2*, Io has a final well-being of $80 + d$ and an expected well-being of $60 + d/2$; this is the EDE of this outcome. The expected value of B is the probability-weighted sum of these EDE values.

How does the expected EDE of B compare to the EDE of A? Let us look at it for each component of our hybrid currency. Prospect B's value in terms of expected well-being is $60 + d/2$, which exceeds the expected well-being yielded by A. Prospect B's expected value in terms of final well-

being is also greater. Our egalitarian view therefore mandates the choice of B.

More generally, in interpersonal trade-offs with inequality, this form of egalitarianism favors the less well off, thereby respecting the separateness of persons. But in intrapersonal trade-offs without inequality, it holds that prudential and moral evaluation converge. It therefore fully respects the unity of the person in such cases.

Nonetheless, it is open to the following objection:

On this view, what one ought to do is, implausibly, affected by the well-being levels of people who have nothing at stake. To illustrate, compare the Intrapersonal, One Person Case (table 5) with the Intrapersonal Scenario (table 2). In the former, this egalitarian view mandates B for every positive d . But in the latter, for some small, positive d , it will prohibit choosing B, because this prospect will generate inequality. This is so even though these prospects affect the well-being of the only people who have something at stake (respectively, Io and Chelsey) in an identical manner. The difference in these verdicts is wholly due to Dave's appearance in the Intrapersonal Scenario, where his presence generates offending inequality. But Dave's well-being is unaffected by the choice between A and B. Indeed, we can imagine that Dave died some time ago and that nothing can disturb the profound security of the well-being he enjoyed. The contrast between these cases therefore highlights the fact that egalitarianism makes what one ought to do depend on the fate of unaffected individuals. This is implausible. Respect for the separateness of persons requires that individuals whose well-being is unaffected by one's choices do not figure in one's decision making; since their well-being is not affected, they have no claim in favor of any of the alternatives. Because, on prioritarian views, the well-being of unaffected parties cannot determine what one ought to do, prioritarian views better respect the separateness of persons, in this sense.²³

In reply: we reject the assumption that the separateness of persons requires that one consider only individuals whose well-being is at stake. The unfairness which concerns egalitarians is essentially about how well some individuals' lives go in comparison to how well other, separate individuals' lives go. A situation in which, due to brute luck, Arnaldur's lifetime quality of life is okay and Bea's is very good is unfair. By contrast, a situation in which, due to brute luck, for one part of his life, Arnaldur

23. We are grateful to a referee for suggesting this objection. For further discussion, see Fleurbaey, "Assessing Risky Social Situations," 665–69; and Shlomi Segall, "Incas and Aliens: The Truth in Telic Egalitarianism," *Economics and Philosophy* 32 (2016): 1–19.

has a merely okay quality of life and for another part of his life, he has a very good quality of life is not unfair in this way. Only in the former situation can Arnaldur legitimately complain of being unfairly disadvantaged vis-à-vis another. He has this complaint even when his well-being cannot be improved and even when the better off Bea exists at a later point in time. Suppose, for example, that Arnaldur develops a debilitating physical illness about which nothing can be done. Bea will also develop this illness, but because she belongs to a future generation, she will have access to a complete cure. Arnaldur can legitimately complain that this differential access to a cure is one way in which “life is unfair.” As a distributor, one should take into account such unfairness when evaluating alternatives. In so doing, one rightly considers the fate of individuals whose well-being is not at stake.

We have argued that for risky, fixed-identity cases, hybrid egalitarianism is superior to hybrid prioritarianism. However, an assessment of these rival views should also take into account how they fare in risky, fixed-number, nonidentity cases. We therefore now take up the task of developing the most plausible versions of prioritarianism and egalitarianism for such cases.

IV

For a possible person who has some chance of never existing, how, if at all, should we value the state in which she never exists? In this section, we consider the answer to this question offered by what we will call the Prospects of Existence View. We argue that the view is coherent, has an interesting rationale, and arrives at the right answer in our opening Interpersonal, One Future Person Case (table 1). But we also argue that it should nonetheless be rejected.

The Prospects of Existence View endorses the following principles for evaluating the prospects of a possible person:

- (i) A state of the world in which this possible person has a positive level of well-being is better for her than a state in which she never exists;
- (ii) A state of the world in which she has negative well-being is worse for her than never existing; and
- (iii) A state of the world in which she never exists is equivalent to her living with a well-being level of zero (at which she would, if rational and considering her self-interest alone, be indifferent between existing and never existing).

It may appear that it makes no sense to value a possible person’s prospects in this manner. To see why, take a person who enjoys positive well-being. If being is better for her than nothingness, then it might seem to

follow that it would have been worse for her if she had never existed. But the latter is regarded by some as an absurdity—if she hadn't existed, nothing would have been better or worse for her, since she would not have been at all.²⁴

However, this allegedly absurd conclusion does not follow. As Gustaf Arrhenius and Wlodek Rabinowicz argue, a triadic relation between a person, a state in which this person leads a good life, and a state in which she never exists holds only if all three relata exist. The states in question are abstract objects and can therefore exist even if they do not actually obtain. But if, as Arrhenius and Rabinowicz assume, a person is a concrete object, then the relation cannot hold if she never exists, since in that case one of the relata does not exist. Therefore, it does not follow from the fact that it is better for her to exist than never to exist that never existing would have been worse for her, since in the latter case one of the relata would have been absent.²⁵

This reasoning establishes that it is coherent to appeal to an existing person's preferences in order to judge that her existence is better for her than her never existing. We submit that it is also coherent to appeal to how good a state would be for a person if she were to come into existence in order to now, in anticipation, comparatively evaluate both states in which she exists and states in which she never exists. To motivate this, reflect first on how one would regard the possibility of a person coming into existence with a negative quality of life. Considering only her interests, one should regard this as worse than her never coming into existence—in line with the way she would, given our assumptions about the measure of well-being, rationally value it if she were to come into existence. Next, reflect on the possibility of a person coming into existence with a positive quality of life. Considering only her interests, one has reason to regard this existence as good—after all, she would, if she were to come into existence, rationally regard her existence as better for herself than never existing.²⁶

The proposed rationale for the Prospects of Existence View, then, is this. Since life at a positive level of well-being is a good and would be better for the person than never existing, the distribution of chances of receiving this good matters morally.²⁷

24. See John Broome, *Ethics out of Economics* (Cambridge: Cambridge University Press, 1999), 186.

25. Gustaf Arrhenius and Wlodek Rabinowicz, "The Value of Existence," in *The Oxford Handbook of Value Theory*, ed. I. Hirose and J. Olson (Oxford: Oxford University Press, 2015), 424–44, 428. For discussion, see Marc Fleurbaey and Alex Voorhoeve, "On the Personal and Social Value of Existence," in *Weighing and Reasoning: Themes from the Work of John Broome*, ed. I. Hirose and A. Reisner (Oxford: Oxford University Press, 2015), 94–109.

26. Here, we follow Arrhenius and Rabinowicz, "The Value of Existence," 430–31. See also Adler, *Well-being and Fair Distribution*, 219–20.

27. We thank a referee for this way of putting the rationale.

Given the proposed equivalence of existing at well-being level zero and never existing, this view adopts the following general principle for our cases: *ceteris paribus*, one ought to treat a case in which a possible person doesn't exist in a state of the world just like a case in which this person exists in that state with a well-being level zero. In the Interpersonal, One Future Person Case (table 1), the Prospects of Existence View therefore regards A as equivalent to a prospect in which either Ann ends up with a final well-being of 60 and Bob with 0 or Ann ends up with 0 and Bob with 60, with both outcomes being equally likely. It regards B as equivalent to a prospect in which either Ann ends up with a final well-being of 40 and Bob with 0 or Ann ends up with 0 and Bob with 80.

In this case, a prioritarian who accepts the Prospects of Existence View will choose A, since: (a) this prospect improves Ann's expectations from the equivalent of an expected well-being of 20 to the equivalent of an expected well-being of 30, which is more important than improving Bob's expectations from the equivalent of an expected well-being of 30 to 40; and (b) prospect A ensures that whoever will exist will have a final well-being of 60, which, for a prioritarian has higher expected value than a fifty-fifty gamble of this person existing at 40 or 80.

An egalitarian who adopts the Prospects of Existence View will also choose A, because it equalizes the value of Ann's and Bob's prospects of existence without a loss in expected total well-being. By being concerned with equality in people's valuable chances of coming into existence, this approach therefore resolves the challenge to egalitarianism raised in the introduction.

However, the Prospects of Existence View yields problematic verdicts in other cases. Consider, for example, the Variable Chance of Existence Case outlined in table 6. In this case, either, with probability p , Jane will exist with a well-being level of $70 + d$, or, with probability $1 - p$, Kathy will exist with a well-being of 70, where $d \geq 0$. You must choose the value of p .

We submit that if Jane's quality of life, were she to exist, would be as good as Kathy's would be if she were to exist instead ($d = 0$), then it is a matter of indifference which probability you choose. Moreover, if Jane

TABLE 6

FINAL WELL-BEING IN THE VARIABLE CHANCE OF EXISTENCE CASE

p		$1 - p$	
Jane	Kathy	Jane	Kathy
$70 + d$			70

would have a better life than Kathy would have ($d > 0$), then it would be best if you maximized the chance that Jane would come into existence and chose $p = 1$.

Neither a prioritarian nor an egalitarian who adheres to the Prospects of Existence View can accommodate these judgments. Suppose first that Jane's quality of life, if she were to come into existence, would be as good as Kathy's would be if she were to exist instead ($d = 0$). From the perspective of the final well-being components of prioritarian and egalitarian views, it is then a matter of indifference which p you choose. However, from the perspective of the value of each person's prospects, you ought to choose p equal to one-half. Since, on the proposed approach, never existing is valued at 0, whoever is less likely to come into existence will be the person with the least valuable prospects. By raising this person's probability of existing, you can improve her prospects to the same extent as you worsen the prospects of the person with the most valuable prospects. Both hybrid prioritarianism and hybrid egalitarianism will therefore require that you equalize Jane's and Kathy's probabilities of existence when $d = 0$.

Now assume that Jane would have a somewhat better life than Kathy would have (d is positive, but small). On the Prospects of Existence View, both hybrid prioritarianism and hybrid egalitarianism can then require that you nonetheless give Kathy a chance of coming into existence. To see why, suppose that you give Kathy no such chance. She then has much worse prospects than Jane: Kathy's will be equivalent to an expected well-being of 0; Jane's will be equivalent to an expected well-being of $70 + d$. By raising Kathy's probability of coming into existence, you improve her prospects from this very low level at a cost of a roughly equally large reduction in Jane's prospects from this much higher level. If one gives extra weight to improving the prospects of the person with worse prospects, for some sufficiently small, positive d , the optimal value for p will be less than unity. Indeed, there is good reason to think that the recommended value of p will be close to one-half. For the Prospects of Existence View is committed to treating this case just like a case in which either, with probability p , Jane has a well-being of $70 + d$ (with d small and positive) and Kathy exists with a well-being level 0, or, with probability $1 - p$, Jane exists with a well-being level 0 and Kathy has a well-being of 70. In the latter case, it is intuitively clear that one ought to give Jane and Kathy a (roughly) equal chance. After all, if these individuals would exist for sure, who would deny Kathy such a (roughly) equal chance at this roughly equally great benefit? But whereas requiring roughly equal chances at a good life would be plausible if Jane and Kathy were to exist for sure, such a requirement is implausible in the Variable Chance of Existence Case. In that case, it is clearly better to ensure the existence of whoever would have a higher level of well-being.

The general problem with the Prospects of Existence View is this. In treating a scenario in which a merely possible individual does not come into existence as akin to a scenario in which this person exists with a zero level of well-being, it ignores the fact that a person who does not come into existence cannot have been wronged. Only people who exist at some time can be wronged by not having their interests properly promoted.²⁸ In the Variable Chance of Existence Case (table 6), if one chooses to create Jane for sure ($p = 1$), then Kathy will never be wronged, because she will never be. This contrasts sharply with a scenario in which Kathy will exist for sure alongside Jane, but either, with probability p , has well-being level 0, or, with probability $1 - p$, has a well-being level of 70. In the latter case, if one chooses $p = 1$, Kathy can rightly object that she was condemned to a life no better than nothingness when she could have instead been given a substantial chance at a good life. By contrast, in the Variable Chance of Existence Case, she cannot make such a complaint if one chooses $p = 1$, because she then doesn't exist. Because this way of being insensitive to the presence or absence of a person who is wronged is problematic, we reject the idea underlying the Prospects of Existence View, namely, that because life would be good for a person, we should care about the distribution of chances at receiving this good. In the next section, we develop a view that does not rest on this erroneous foundation.

V

On what we will call the Conditional on Existence View, in evaluating a potential outcome of an alternative, one adopts the following principles:

- (i) One should consider only the individuals who exist in that outcome;
- (ii) Of these individuals, one should be concerned with
 - (ii-a) their final well-being; and
 - (ii-b) their expected well-being conditional on their existence.

On this view, while we should not care about people's chances of coming into existence (because they have no complaint if they do not come into existence), we should care about how they would (or might) fare if they were to exist. To illustrate, consider again the Variable Chance of Existence Case (table 6) in a version in which Jane's well-being would be greater than Kathy's would be ($d > 0$). In evaluating different al-

28. Here, we are agreeing with Melinda Roberts, *Child versus Child-maker: Future Persons and Present Duties in Ethics and the Law* (Lanham, MD: Rowman & Littlefield, 1998); and Peter Vallentyne, "Review of *Child versus Child-maker: Future Persons and Present Duties in Ethics and the Law* by Melinda Roberts," *Noûs* 34 (2000): 634–47.

ternatives (which in this case, recall, are different chances of existence for Jane and Kathy), one has two possible outcomes to consider. In one outcome (which occurs with probability p) Jane exists with a final well-being of $70 + d$ and an expected well-being conditional on her existence of $70 + d$. (If she exists, she invariably has this level of well-being. Her expected well-being conditional on her existence is therefore identical to her final well-being in this simple case.) In the other outcome (which occurs with probability $1 - p$), Kathy exists with a final well-being of 70 and a conditional expected well-being of 70. Obviously, since $d > 0$, one maximizes expected value by ensuring Jane's existence ($p = 1$). (This is true whether one is a prioritarian or an egalitarian.) This approach therefore arrives at the right answer in this case.

Next, consider again the Interpersonal, One Future Person Case (table 1). In this case, a prioritarian who adopts the Conditional on Existence View is concerned with a hybrid of final and conditional expected well-being. Such a form of prioritarianism mandates A, since it regards it as more important to improve Ann's final and conditional expected well-being from 40 to 60 than to improve Bob's from 60 to 80.

Matters are less straightforward, however, for an egalitarian who follows the Conditional on Existence View. In the Interpersonal, One Future Person Case, this form of egalitarianism considers only Ann's situation in *S1* and only Bob's situation in *S2*. It does not evaluate Ann's situation in comparison with Bob's. This absence of comparative evaluation reflects the absence of the kinds of inequality that concern egalitarians. Since Ann and Bob will not both exist, if you choose B, Ann will not be worse off than Bob; nor will Bob be worse off than Ann. The familiar source of unfairness therefore does not exist in this instance. It follows that the fairness-based justification for equality in (conditional) expected well-being among actual individuals which we proposed in Section III—that such equality of chances can mitigate the unfairness of unequal final well-being—cannot be invoked here either. It is therefore in question whether the proposed form of egalitarianism can yield the desired verdict that we should choose A.

To ensure that it can, we need to address an issue we have so far left undiscussed. In every outcome, a person has a particular final well-being and had (at the moment of decision) a particular conditional expected well-being. This gives rise to the following question: How do the two elements of the currency of distributive ethics—final well-being and conditional expected well-being—jointly determine how much a person has received of this currency in a particular outcome? From our stipulation (in the introduction) that a person's expected holdings of this currency track his prudential interest, it follows that when this currency is a mix of final and conditional expected well-being: (a) an increment in final well-being always adds the same amount to the currency a person

has received; and (b) an increase in conditional expected well-being always increases the amount of the currency received. While (a) specifies that the contribution of an increment in final well-being to the currency of distributive ethics is both positive and constant, (b) says only that the marginal contribution of conditional expected well-being is positive. Our definition of the currency of distributive ethics therefore leaves unspecified whether the latter's marginal contribution is increasing, constant, or decreasing. Egalitarianism is also silent about this issue—it says only that one should care about inequalities in this hybrid currency but says nothing about how, precisely, conditional expected well-being contributes to this currency. Pluralist egalitarians can use this degree of freedom to specify a way in which increments in conditional expected well-being contribute to the hybrid currency that yields the right answer in the Interpersonal, One Future Person Case (table 1) and in other cases. The following specification of (b) yields the desired result: the marginal contribution of conditional expected well-being to the currency of egalitarian ethics is positive, but decreasing. In other words, the higher an individual's conditional expected well-being, the less a given increment in his conditional expected well-being adds to the currency he receives. Having very good rather than moderately good conditional prospects generates less of an increment in this currency than having moderately good rather than merely okay conditional prospects.²⁹

If it incorporates this idea, then our pluralist egalitarian view holds that in the Interpersonal, One Future Person Case (table 1), the value of improving Ann's conditional expected well-being from 40 to 60 exceeds the value of improving Bob's conditional expected well-being from 60 to 80, because the former improvement takes place from a lower level. Prospect A is therefore preferred.

29. It may be helpful to put the preceding in slightly more formal terms. In every outcome (which is determined by the choice of a prospect and the revelation of a state of nature), a person i has final well-being w_i and had (at the moment of decision) conditional expected well-being cew_i . The question is how these two combine to form a single measure m_i of the amount of the hybrid currency that this person has received in a given outcome. We are proposing that an egalitarian adopt the following measure: $m_i = w_i + \varphi(cew_i)$, in which φ is concave to reflect the diminishing marginal contribution of conditional expected well-being to this currency.

This "diminishing marginal contribution" principle bears a family resemblance to prioritarianism for conditional expected well-being. Are we therefore not simply proposing a form of pluralist egalitarianism which gives some weight to an element of conditional expected well-being prioritarianism? No. Conditional expected well-being prioritarianism is a combination of a claim about the currency of distributive ethics (that it consists in conditional expected well-being) with a claim about how this currency should be distributed (in a prioritarian fashion). By contrast, our "diminishing marginal contribution" principle is about how one of the elements of the currency of ethics (conditional expected well-being) contributes to how much of that currency an individual receives.

In sum, our pluralist egalitarian view arrives at the right conclusion in our opening case if it relies on this “diminishing marginal contribution of conditional expected well-being to the currency of ethics” principle. But one must ask: is this principle defensible? Isn’t an appeal by egalitarians to such a nonegalitarian principle ad hoc?

In reply: our defense of this principle appeals both to its implications in cases and to its fit with still deeper principles. The principle gains some support from the fact that it enables an otherwise appealing egalitarian view to arrive at the right answer in the Interpersonal, One Future Person Case (table 1). As we have seen, it does so not by limiting this egalitarian view, but rather by filling in a blank space in this view. Moreover, it is not an ad hoc, but rather a natural addition to this pluralist, hybrid egalitarian view. For it can be justified by an appeal to the same foundational moral ideal that justifies the other elements of this pluralist view: respect for the difference between intra- and interpersonal trade-offs. To see why, compare the Interpersonal, One Future Person Case (table 1) with a variant of our Intrapersonal, One Person Case (table 5) in which $d = 0$, so that, in the latter case, both A and B offer equal conditional expected well-being. Under this assumption, both cases have strong similarities. In both cases, only one person will exist. In both cases, if one chooses A, the person who exists will have a final and conditional expected well-being of 60. Moreover, in both cases, if one chooses B, in state of the world $S1$, whoever exists will have a final well-being of 40; and in $S2$, whoever exists will have a final well-being of 80. The only difference is that, if one chooses B in this variant of the Intrapersonal, One Person Case, the person who exists will have a conditional expected well-being of 60 because both potential futures are hers, whereas in the Interpersonal, One Future Person Case, the person who exists will either have a conditional expected well-being of 40 or a conditional expected well-being of 80 because these potential futures belong to different people. According to the diminishing marginal contribution principle, this difference is significant. Prospect B is permissible in the Intrapersonal, One Person Case because it maximizes the conditional expected well-being of the only person whose well-being is at issue. In contrast, B is impermissible in the Interpersonal, One Future Person Case because it is more important to raise one possible person’s conditional expected well-being from 40 to 60 than to raise a different possible person’s conditional expected well-being from 60 to 80. This principle therefore marks the difference between intrapersonal trade-offs without inequality (in which it endorses maximizing expected well-being) and interpersonal trade-offs without inequality (in which it gives priority to improving the prospects of the possible person with the worst conditional prospects).

Every element of our proposed pluralist, hybrid egalitarian view therefore has the same deep rationale: respect for both the unity of the individual and the separateness of persons. The unity of the individual gives us reason to maximize a person's expected well-being. The separateness of persons gives us reason to care about unfair inequality in conditional expected well-being and in final well-being. It also gives us reason to improve the prospects of those with lower conditional expected well-being even in the absence of unfair inequality.

VI

The proposed pluralist, hybrid egalitarian view arrives at the right conclusion in all cases so far surveyed. In this, it is superior to hybrid prioritarianism, which, as the Intrapersonal, One Person Case (table 5) demonstrates, fails to fully respect the unity of the individual. In closing, we present one further advantage that this form of egalitarianism has over prioritarianism.

Consider the Costly Equality Case outlined in table 7, in which $20 > d > 0$.

You can either (prospect A) moderately benefit both Leo and Maria, ensuring they will lead equally good lives, or (prospect B) provide a smaller benefit to Leo and an especially large benefit to Maria. By choosing A, you achieve equality in a manner that improves the well-being of the worst off. However, doing so is costly in terms of total final and expected well-being. The variable d represents this cost. Ask yourself for which d you ought to be indifferent between A and B in this case. This "indifference d " is the largest cost you should be willing to accept to improve the situation of the worst off.

Now consider the Costly Improvements in Worst Prospects Case outlined in table 8, in which, again, $20 > d > 0$.

In this case, either Nona or Oliver will come into existence. Both are equally likely to exist. You can either (prospect A) arrange matters so that whoever will exist receives a moderate benefit, or (prospect B) arrange matters so that if Nona comes into existence, she will receive a

TABLE 7

FINAL WELL-BEING IN THE COSTLY EQUALITY CASE

	Leo	Maria
Prospect A	$60 - d$	$60 - d$
Prospect B	40	80

TABLE 8
 FINAL WELL-BEING FOR THE COSTLY IMPROVEMENT
 IN WORST PROSPECTS CASE

	S1 (0.5)		S2 (0.5)	
	Nona	Oliver	Nona	Oliver
Prospect A	60 - d		60 - d	
Prospect B	40		80	

smaller benefit and if Oliver comes into existence, he will receive an especially large benefit. By choosing A, you improve the fate of the possible person with the worst conditional expected well-being, but doing so is costly in terms of expected total final and conditional expected well-being. The variable d represents this cost. Ask yourself for which d you ought to be indifferent between A and B in this case. This “indifference d ” is the largest cost you should be willing to accept to improve the situation of the possible person with the lowest conditional expectations.

Should your “indifference d ” in the Costly Equality Case (table 7) be the same as your “indifference d ” in the Costly Improvement in Worst Prospects Case (table 8)? We think not. In the former, you can prevent inequality between persons who will both exist. In the latter, there will be no such inequality. We therefore believe that you ought to accept a larger cost to raise Leo to the same level as the coexistent Maria than you ought to be willing to accept to raise Nona’s well-being, if she were to exist, to the same level that Oliver would be at if he were to exist instead of her. In deciding whether to improve the prospects of a possible person with less good prospects at a cost to a possible person with better prospects, you should allow a larger cost when these possible people will both be actual.

Our egalitarian view concurs. In the Costly Equality Case (table 7), it holds that improving the lot of the worst off is important because it reduces inequality in the distribution of final and expected well-being. In the Costly Improvement in Worst Prospects Case (table 8), it registers no such egalitarian reason to favor the worst off. It registers only the diminishing marginal contribution of conditional expected well-being to the distributive currency. It will therefore tolerate a larger cost in the former than in the latter.

By contrast, prioritarian views will not accept a costlier transfer in the former case. On these views, improving someone’s final and conditional expected well-being from 40 to 60 - d leads to the same gain in moral value whether or not there is another, better-off individual, while

having someone's final and conditional expected well-being be $60 - d$ rather than 80 leads to the same loss in moral value, independently of the existence of another, worse off individual. The point at which this gain and loss are equal is therefore the same in both cases. In this pair of cases, prioritarianism therefore displays an unappealing insensitivity to the presence or absence of inequality between people who exist.

VII

We have examined both familiar, risky cases in which people's identities are fixed and less familiar cases in which our choices influence the prospects and/or identities of the people who will exist, but not the number of people who exist. Our abstract cases model considerations that are relevant when we must make risky choices to improve the well-being of existing people. They also represent considerations that are relevant when we must try to benefit future people and we have only probabilistic information about who these people will be. (Real-world examples of the latter are aid to future children of families in need, or policies that help future generations adapt to climate change.) Our central conclusions are the following.

First, the currency of well-being-based distributive ethics is a combination of an individual's final well-being and his expected well-being conditional on his existence.

Second, there is a pluralist egalitarian view, based on the equally distributed equivalent, which: (a) fully respects both the unity of the individual and the separateness of persons and (b) is rightly more willing to favor a costly improvement in the expectations of the person with the worst prospects when this reduces inequality between actual people than when it does not reduce such inequality. Even the most plausible prioritarian view fails in these respects.

In sum, for all people, actual and possible, the currency of distributive ethics is a combination of final well-being and conditional expected well-being. One should distribute this currency in accordance with egalitarian, rather than prioritarian, principles.

30. For a related claim that prioritarianism fails to recognize the importance of inequality in risky, fixed-identity cases, see John Broome, "Equality versus Priority: A Useful Distinction," *Economics and Philosophy* 31 (2015): 219–28.