

Priority, Not Equality, for Possible People*

Jacob M. Nebel

How should we choose between uncertain prospects in which different possible people might exist at different levels of well-being? Alex Voorhoeve and Marc Fleurbaey offer an egalitarian answer to this question. I give some reasons to reject their answer and then sketch an alternative, which I call *person-affecting prioritarianism*.

How should we choose between uncertain prospects in which different possible people might exist at different levels of well-being? Alex Voorhoeve and Marc Fleurbaey offer an egalitarian answer to this question.¹ I explain their motivation for this answer in Section I. In Sections II and III, I give some objections to their version of egalitarianism. In Section IV, I sketch an alternative account of their central intuition. This account, which I call *person-affecting prioritarianism*, avoids my objections to Voorhoeve and Fleurbaey's egalitarianism and many of the objections to other versions of prioritarianism.

I. HYBRID EGALITARIANISM

Suppose that you must choose between two uncertain prospects, *A* and *B*, in which one person will come into existence. You are uncertain not only of the future person's welfare but also of her identity, as depicted in table 1.² Either Ann or Bob will exist, but you don't know which. In state 1, Ann exists; in state 2, Bob exists. States 1 and 2 are equally probable on

* For helpful comments and discussion, I am grateful to Kara Dreher, Marc Fleurbaey, Robert Long, Michael Otsuka, Ketan Ramakrishnan, Samuel Scheffler, Trevor Teitel, and the editors.

1. Alex Voorhoeve and Marc Fleurbaey, "Priority or Equality for Possible People?," *Ethics* 126 (2016): 929–54.

2. I have slightly modified the case discussed by Voorhoeve and Fleurbaey. The case is originally due to Michael Otsuka, "Prioritarianism and the Separateness of Persons," *Utilitas* 24 (2012): 365–80, 369.

TABLE 1

RISKY NONIDENTITY CASE

PROSPECT	STATE 1 (0.5)		STATE 2 (0.5)	
	Ann	Bob	Ann	Bob
<i>A</i>	49			49
<i>B</i>	0			98 + <i>d</i>

your evidence. The number in a person's cell (if there is one) represents her quantity of well-being on an interpersonal ratio scale—that is, a scale on which ratios between numbers reflect how many times better off one person is than another. (An empty cell represents nonexistence.) I assume, following Voorhoeve and Fleurbaey, that our measure of well-being is derived from an idealized preference relation satisfying the axioms of expected utility theory.³ On this view, a prospect offers you greater expected well-being just in case a fully rational agent concerned solely with your interests would prefer it for your sake. And if such an agent would prefer something for your sake, then it is better for you.⁴

In table 1, it seems harder to justify the risky prospect *B* than it would be if both futures were Ann's. For suppose that you choose *B* and state 1 obtains. If both futures had been Ann's, then (for any $d > 0$) you could have justified your choice to her on the grounds that *B* was in her expected interests—it would have maximized her expected well-being—given what you knew at the time. But this justification is not available in the risky nonidentity case. *B* does not maximize Ann's expected well-being, since the future in state 2 would be Bob's. Choosing *B* would seem to sacrifice Ann's interests for the sake of Bob's. It therefore seems that, for some $d > 0$ (and, of course, all $d \leq 0$), you should choose *A* to ensure that Ann is not so badly off.

This judgment could perhaps be explained by a *prioritarian* theory of distribution. Prioritarians believe, in Derek Parfit's words, that "we have stronger reasons to benefit people the worse off these people are."⁵ *Outcome* prioritarians believe that such reasons are stronger because benefit-

3. Expected utility theory provides a cardinal scale. A ratio scale requires a meaningful zero level. For possible ways of defining such a level, see Charles Blackorby, Walter Bossert, and David Donaldson, *Population Issues in Social Choice Theory, Welfare Economics, and Ethics* (Cambridge: Cambridge University Press, 2005), chap. 2.

4. This does not entail that preference satisfaction is what makes our lives go well, or that goodness should be reduced to or analyzed in terms of rational preferences.

5. Derek Parfit, "Another Defence of the Priority View," *Utilitas* 24 (2012): 399–440, 401.

ing the worse-off would bring about a better outcome, since (on their view) increases in well-being make a diminishing marginal contribution to the goodness of outcomes. Outcome prioritarrians would claim that, for some $d > 0$ in table 1, we ought to choose *A*, because doing so would, in expectation, bring about a greater sum of *priority-weighted* well-being.⁶

Outcome prioritarianism, however, registers no difference between the risky nonidentity case, in which *B* sacrifices Ann's interests for the sake of another person, and a purely intrapersonal version of the case, in which both futures would be Ann's. In such a purely intrapersonal case, *B* would run the risk of making Ann worse off, but this risk could be justified to Ann as being in her expected interests. Given our measure of well-being, *B* would be preferred for Ann's sake if both futures would be hers, because it would offer her greater expected well-being. The sum of priority-weighted well-being, however, does not depend on which people exist. Outcome prioritarianism would recommend *A* for the same values of d in both the risky nonidentity case and its intrapersonal analogue. Outcome prioritarianism, therefore, fails to respect the morally significant difference between intrapersonal and interpersonal trade-offs.⁷

How, then, can we explain the judgment about the risky nonidentity case in table 1? Voorhoeve and Fleurbaey develop a version of egalitarianism which secures this judgment. This may seem surprising, because there is no objectionable inequality in the risky nonidentity case: it is not unfair that, under *B*, Ann would be worse off than Bob would have been, had he existed. The only objectionable inequalities are between actual, not merely possible, people.⁸ But Voorhoeve and Fleurbaey propose a way of capturing the intuitive judgment without appealing to inequalities between merely possible people.

The relevant aspect of Voorhoeve and Fleurbaey's egalitarianism is a novel view about the "currency" of distributive ethics—that is, the thing whose distribution we ought to care about being patterned in a certain way, and the metric by which other goods ought to be distributed.⁹ Voorhoeve and Fleurbaey's proposed currency is a hybrid of two ingredients. The first is a person's final well-being—that is, how well off she ends up in each outcome. The second is a person's expected well-being conditional on her existence—that is, how well off she would be, in expectation, con-

6. A person's priority-weighted well-being in an outcome is given by some strictly increasing, strictly concave function of her well-being.

7. Otsuka, "Prioritarianism and the Separateness of Persons," 369; see also Michael Otsuka and Alex Voorhoeve, "Why It Matters That Some Are Worse Off Than Others: An Argument against the Priority View," *Philosophy and Public Affairs* 37 (2009): 171–99.

8. Otsuka, "Prioritarianism and the Separateness of Persons," 370.

9. G. A. Cohen, "On the Currency of Egalitarian Justice," *Ethics* 99 (1989): 906–44; see also Amartya Sen, "Equality of What?," in *The Tanner Lectures on Human Values*, ed. Sterling McMurrin (Cambridge: Cambridge University Press, 1980), 1:196–220.

TABLE 2

HYBRID HOLDINGS IN THE RISKY NONIDENTITY CASE

PROSPECT	STATE 1 (0.5)		STATE 2 (0.5)		EXPECTED VALUE
	Ann	Bob	Ann	Bob	
A	$49 + \sqrt{49} = 56$			$49 + \sqrt{49} = 56$	56
B	0			$100 + \sqrt{100} = 110$	55

sidering only the outcomes in which she exists. Importantly, however, although a person’s final well-being makes a constant marginal contribution to her holdings of the hybrid currency, her conditional expected well-being makes a decreasing marginal contribution to that currency. To make this aspect of the view concrete, I assume that one’s holdings of the hybrid currency can be represented by the sum of one’s final well-being and the square root of one’s conditional expected well-being.¹⁰ The value of each outcome can then, according to Voorhoeve and Fleurbaey’s view (which I shall call *hybrid egalitarianism*), be determined by applying a standard egalitarian principle to the distribution of hybrid currency. The value of a prospect is then identified with its expected value.

To see how hybrid egalitarianism applies to the risky nonidentity case in table 1, suppose that $d = 2$, so that Bob’s well-being in *B* comes to 100 ($= 98 + 2$). Table 2 shows the resulting distributions of hybrid currency, as well as the expected value of each prospect according to hybrid egalitarianism.

Since these outcomes involve no inequality, the value of each outcome can be equated with the holdings of whichever person exists in that outcome. The expected value of each prospect is just the sum of the values of its outcomes, weighted by their probabilities. That is why, in this setup, *A* is better than *B*. But, unlike outcome prioritarianism, hybrid egalitarianism registers a morally significant difference between interpersonal and intrapersonal trade-offs. For if both futures were Ann’s, then *B* would maximize her expected holdings of the hybrid currency (for any positive d), since it would increase her expectation of both of its ingredients.¹¹

10. Voorhoeve and Fleurbaey do not mention the square root function in particular—only that it must be concave (“Priority or Equality,” 950 n. 29). I choose this function because it simplifies the calculations below. It does not affect the substance of my arguments.

11. For any positive d , $0.5(0 + \sqrt{0.5(0) + 0.5(98 + d)}) + 0.5(98 + d + \sqrt{0.5(0) + 0.5(98 + d)}) > 0.5(49 + \sqrt{0.5(49) + 0.5(49)}) + 0.5(49 + \sqrt{0.5(49) + 0.5(49)})$.

II. THE HYBRID CURRENCY AND PERSONAL GOOD

Let me offer a different kind of risky nonidentity case, depicted in table 3. Suppose that you are a doctor. Some patient of yours will have a child. The patient might, with probability p , have a rare medical condition. If she has this condition and you do nothing, then she will have a child (call her Cat) whose life would be barely worth living. If your patient lacks the condition and you do nothing, then she will have a different child (call him Dan), whose life would be excellent.

You cannot treat the condition. But you can intervene with the patient's uterine environment so as to settle which zygote will be produced, regardless of the condition. Doing so, however, would make the resulting life *worse*, no matter what. Table 3 represents the options available to you.

In state 1, the patient has the condition; in state 2, she lacks it. A is the prospect of doing nothing. B is the prospect of intervening so as to cause Cat to exist. C is the prospect of intervening so as to cause Dan to exist.

Prospect A would bring about a better life, regardless of whether your patient has the condition. It seems clear to me that A is best, and that you therefore ought to do nothing, for any $1 \geq p \geq 0$ and any $d > 0$.

According to hybrid egalitarianism, however, for any probability $1 > p > 0$, there is some $d > 0$ small enough that it would be wrong to choose A . Suppose, for example, that $p = 0.55$ and $d = 1$. The resulting distributions of final well-being, conditional expected well-being, and hybrid currency are shown in table 4. (I omit C because its figures are the same as B 's, but for Dan.)

Because there is no inequality in these outcomes, the value of each outcome is just the amount of hybrid currency held by each person who exists in that outcome. A 's expected value is, therefore, 41.6 ($= 0.55(2) + 0.45(90)$). B 's is 42 ($= 0.55(6) + 0.45(86)$). So, in this setup, hybrid egalitarianism ranks B better than A . Hybrid egalitarianism yields this verdict because conditional expected well-being makes a diminishing marginal

TABLE 3
THE ZYGOTIC SELECTION CASE

PROSPECT	STATE 1 (p)		STATE 2 ($1 - p$)	
	Cat	Dan	Cat	Dan
A	1			81
B	$1 - d$		$81 - d$	
C		$1 - d$		$81 - d$

TABLE 4

FINAL, EXPECTED, AND HYBRID DISTRIBUTIONS
IN THE ZYGOTIC SELECTION CASE

DISTRIBUTION	STATE 1 (0.55)		STATE 2 (0.45)	
	Cat	Dan	Cat	Dan
A:				
Final	1			81
Expected	1			81
Hybrid	$1 + \sqrt{1} = 2$			$81 + \sqrt{81} = 90$
B:				
Final	0		80	
Expected	$0.55(0) + 0.45(80) = 36$		$0.55(0) + 0.45(80) = 36$	
Hybrid	$0 + \sqrt{36} = 6$		$80 + \sqrt{36} = 86$	

contribution to the hybrid currency. If conditional expected well-being instead made a constant marginal contribution to the hybrid currency, then A would be better, because there would then be no disadvantage to offering different prospects to different possible people.¹²

This verdict seems wrong. For any $1 \geq p \geq 0$ and any $d > 0$, A seems better. We can defend this judgment by appealing to

Statewise Dominance: For any prospects X and Y, if the outcome of X is better than the outcome of Y under every state of the world, then X is better than Y.¹³

Hybrid egalitarians would, I think, accept statewise dominance.¹⁴ They would also agree that A's outcome in state 2 is better than B's, because it contains a greater quantity of hybrid currency (with no inequality). If A's outcome in state 1 is also better than B's, then statewise dominance implies that A is better. And A's outcome in state 1 does seem better, because it makes Cat better off. We can appeal to

12. Suppose, e.g., that one's quantity of the hybrid currency is given by the sum of final and (unweighted) conditional expected well-being. Then A's expected quantity of that currency would be $0.55(1 + 1) + 0.45(81 + 81) = 74$, whereas B's would be $0.55(0 + 36) + 0.45(80 + 36) = 72$.

13. John Quiggin, "Stochastic Dominance in Regret Theory," *Review of Economic Studies* 57 (1990): 503–11.

14. 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. Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler (Oxford: Oxford University Press, 2013), 113–28.

The Weak Principle of Personal Good: For any outcomes X and Y containing the same people, if X is better for everyone than Y , then X is better than Y .¹⁵

This principle implies that A 's outcome in state 1 is better than B 's, because it is better for everyone (i.e., Cat). Since A is better than B under both states 1 and 2, it must be better.

In applying the weak principle of personal good, I have assumed that an outcome is better for someone if she has greater well-being in that outcome. That is why A 's outcome in state 1 is better for Cat than B 's. Hybrid egalitarians, however, might reject this assumption. They might claim that their hybrid currency, rather than well-being, tracks personal good—that is, goodness for a person—so that A 's outcome in state 1 would count as worse for Cat than B 's. But this move would be hard to square with Voorhoeve and Fleurbaey's measure of well-being, which I introduced at the beginning of this article. On this view, quantities of well-being are derived from idealized preferences satisfying the axioms of expected utility theory. A prospect offers you greater expected well-being just in case a fully rational agent concerned solely with your interests would prefer it for your sake—in which case it is better for you.¹⁶ This makes it highly plausible that an outcome in which you have greater well-being is better for you. So, because Cat's well-being is greater in A 's outcome in state 1 than in B 's, A 's outcome is better for her.

Hybrid egalitarians should not jettison Voorhoeve and Fleurbaey's measure of well-being. We must earn the right to represent well-being with numbers. If our numbers are values of an expectational utility function which represents an idealized preference relation, then we know what they mean. But if they are supposed to represent something else—primitive quantities of well-being—then they are far from obviously meaningful.¹⁷ I do not find myself equipped with a pretheoretic ratio scale of well-being. Moreover, Voorhoeve and Fleurbaey's measure of well-being is part of what underlies their objection to outcome prioritarianism: outcome prioritarianism recommends that we act against people's expected interests in purely intrapersonal trade-offs, as we saw in Section I. This recommendation is highly implausible if, given our measure of well-being, any prospect that of

15. See John Broome, *Weighing Goods: Equality, Uncertainty and Time* (Oxford: Wiley-Blackwell, 1991). I only assume the weak principle of personal good for outcomes, not for prospects.

16. The move from preferability to betterness might be questioned. But that would be a nonstarter, since hybrid egalitarianism would still violate the (no less plausible) principle that, for any outcomes X and Y containing the same people, if X would rationally be preferred to Y for the sake of each person, then X is better than Y .

17. Hilary Greaves, "Antiprioritarianism," *Utilitas* 27 (2015): 1–42, however, argues that prioritarians are better off taking quantities of well-being as primitive.

fers you higher expected well-being ought to be preferred for your sake.¹⁸ But it may seem much less implausible if quantities of well-being are understood in some other way.

Suppose that hybrid egalitarians stick with Voorhoeve and Fleurbaey's measure of well-being. Then, in order to reconcile hybrid egalitarianism with the weak principle of personal good, they would have to deny that outcomes in which you have greater well-being are always better for you, while maintaining that prospects in which you have greater expected well-being are always better for you. This move is somewhat ad hoc, but it may have some precedent. Egalitarians sometimes claim that people have a non-welfare-based interest in fair treatment.¹⁹ On this kind of view, you have reason to prefer, for your sake, outcomes brought about by deliberation that treats you fairly, even when such treatment does not contribute to (and may come at the expense of) your well-being. This might seem to support a similar interest in conditional expected well-being: you have reason to prefer, for your sake, outcomes brought about by the choice of prospects which offer you greater expected well-being, conditional on your existence, over and above any contribution made by such choices to your final well-being. Such an interest could make it rational to prefer *B*'s outcome in state 1 for Cat's sake and therefore make this outcome better for Cat. But such an interest in conditional expected well-being cannot be grounded in a concern for fairness. It is not unfair to Cat that her existence in *A* is less robust than it is in *B*. She would have no grounds for complaint if we choose *A*. For if we choose *A* and she exists, then her life is better than it would have been had we chosen *B*, and no worse than anyone else's.

Hybrid egalitarians' interest in conditional expected well-being must, therefore, be independent of our concern for fair treatment. But it does not seem rational to care about conditional expected well-being in its own right—that is, over and above any contribution to our final well-being, and apart from our interest in being treated fairly. On the view we are considering, an outcome in which you have a well-being of *w* would be better for you if it came about via a prospect in which your expected well-being, conditional on your existence, was greater than *w*. This view would seem to imply that if someone has been buying lottery tickets for you (such that you would get the winnings but they pay the cost), things are much better for you than you might have thought, even if you never know it and your tickets never win, and even though these losing tickets make no contribution to your well-being. This view would also seem to im-

18. See Michael Otsuka, "Prioritarianism and the Measure of Utility," *Journal of Political Philosophy* 23 (2015): 1–22.

19. See, e.g., Alex Voorhoeve and Marc Fleurbaey, "Egalitarianism and the Separateness of Persons," *Utilitas* 24 (2012): 381–98, 396.

ply that, no matter how miserable someone's life is, things are ultimately quite good for her so long as her expected well-being was (perhaps at some privileged time, such as conception) sufficiently high. These claims are hard to believe, because these facts about our prospects seem irrelevant to our attitudes about our final situation. We do not regard our circumstances as much more desirable on the supposition that we had good chances, relative to some agent's evidence at some earlier time, of a better life. Even if our knowledge of those chances would have made us happy, the facts known would not have made things better for us. But such facts would be relevant if chances at greater well-being really made outcomes better for us, as they would if the hybrid currency tracked personal good. So, I suspect, the hybrid currency does not track personal good; well-being does.²⁰

An increase in conditional expected well-being is not, in itself, worth wanting at the expense of final well-being. So *A*'s outcome in state 1 is better for Cat than *B*'s. So, by the weak principle of personal good, *A*'s outcome in state 1 is better overall. Since *A* would also be better if state 2 obtains, statewise dominance implies that *A* is better. This means that even if conditional expected well-being is an ingredient of the currency of distributive ethics, its marginal contribution to that currency does not decrease.

III. UNFAIR CHANCES WITHOUT CHANCES OF UNFAIRNESS

But conditional expected well-being is not a proper ingredient of the currency of egalitarian distribution. Suppose, in our zygotic selection case, that you have the same options as in table 3. Let p be such that, when $d = 1$ in that case, *A* would be better than *B* and *C* even if conditional expected well-being had a decreasing marginal impact on our object of distributive concern. For example, let $p = 0.2$.²¹

Now, I have some good news and some bad news about your patient. The good news: she will have twins! The bad news: if she has the rare medical condition (which she probably doesn't), this will affect both twins. Although you can still intervene to fix the identity of one twin to be either Cat or Dan, you can do nothing to affect either the identity or well-being of the other (whom I'll call Eve). Your options are shown in table 5. Again, *A* is the prospect of doing nothing; *B* is the prospect of intervening so as to ensure that Cat exists; *C* is the prospect of intervening so as to ensure that Dan exists.

20. My claim here is not that a chance of a good has no value, or is not worth wanting. My claim is merely that it is not worth wanting over and above its expected contribution to one's well-being. So the unrealized chance of a good cannot make your final situation better for you, unless such a chance contributes to your final well-being.

21. This would make *A*'s expected quantity of hybrid currency $72.4 = 0.2(2) + 0.8(90)$ and *B*'s $72 = 0.2(0 + \sqrt{0.2(0) + 0.8(80)}) + 0.8(80 + \sqrt{0.2(0) + 0.8(80)})$.

TABLE 5
 ZYGOTIC SELECTION CASE WITH TWINS

PROSPECT	STATE 1 (0.2)			STATE 2 (0.8)		
	Cat	Dan	Eve	Cat	Dan	Eve
<i>A</i>	1		1		81	81
<i>B</i>	1 - <i>d</i>		1	81 - <i>d</i>		81
<i>C</i>		1 - <i>d</i>	1		81 - <i>d</i>	81

In the zygotic selection case with twins, you ought to choose *A* for any $d > 0$. This judgment would be shared by most utilitarians, egalitarians, and prioritarians. Fixing the child’s identity would bring about lower total and average well-being, would be worse for the worst-off people in any given outcome, and would introduce inequality in final well-being. And, intuitively, the presence of an unaffected twin cannot make it the case that we ought to fix the child’s identity when we would have otherwise been obligated not to fix it.

According to hybrid egalitarianism, however, it can be wrong to choose *A*, for some $d > 0$. This is because each outcome in *A* involves vast inequalities in conditional expected well-being. Eve’s expected well-being, conditional on her existence, is 65 ($= 0.2(1) + 0.8(81)$), which is much greater than Cat’s and much lower than Dan’s. So if conditional expected well-being is a component of our currency, then *A*’s outcomes may involve a more unequal distribution of that currency than *B*’s and *C*’s. If d is small enough, this inequality can outweigh *A*’s greater quantity of final well-being.

It is not news that egalitarianism makes the value of an outcome depend on the welfare of unaffected individuals. Many people find that implication objectionable. In response, many egalitarians claim that the value of an outcome should depend on the welfare of unaffected individuals, because their welfare can affect the fairness of an outcome. Voorhoeve and Fleurbaey, for example, claim that it is bad, because unfair, that people in previous generations were better or worse off than we are now, even though we cannot now affect their lives.²² But this response would seem to favor *A*, since it is sure to bring about a more equal distribution of final well-being.

What counts against *A*, for hybrid egalitarianism, is the inequality in conditional expected well-being. But it does not seem bad or unfair that Cat’s conditional expected well-being is lower than Eve’s in state 1, or that Eve’s is lower than Dan’s in state 2. An equal distribution of chances seems

22. Voorhoeve and Fleurbaey, “Priority or Equality,” 943.

important because it mitigates the badness of inequality in final well-being. It seems less unfair to make some worse off than others if each person receives an equal chance of being worse off. Indeed, that is how Voorhoeve and Fleurbaey motivate the importance of expected well-being within their egalitarian framework.²³ But it is hard to see why an unequal distribution of chances should be considered unfair when there is no chance of an unequal outcome.

To see why *A* is not an unfair distribution of chances, consider the matter in terms of complaints. In ordinary cases of outcome inequality, it seems reasonable to complain that one is worse off than others through no fault of one's own. This makes it plausibly unfair that one is worse off than others, through no fault of one's own.²⁴ The complaint is stronger, and the unfairness greater, if one was also given a worse chance at being well off than others. But if (like Cat in *A*) one had no chance of being worse off than others, it does not seem reasonable to complain about the fact that one was given a worse chance at being well off than others. For consider what we would have reason to wish or to hope for if Cat's complaint were reasonable. Suppose we choose *A* and that state 1 obtains. If we could make Cat better off, we would of course have reason to do so. But that reason is grounded in our concern for Cat's well-being, not in a concern for equality in conditional expected well-being. To see this, suppose that although we couldn't have made Cat better off, we could have acquired some information that would have made it more probable on our evidence that state 1 would obtain. Acquiring this information would have lowered Eve's conditional expected well-being, thereby decreasing the relevant inequality. But that mere acquisition of information would not, I believe, have made the outcome in any way better. A complaint that would be alleviated by acquiring such information is not a reasonable complaint.

Unequal chances may seem unfair when an agent causes different people to have different chances, thereby benefiting people unequally. For example, if I flip a biased coin to determine the allocation of some good to which two people have equal claims, I cause one of these people to have better prospects than the other. Arguably, this is unfair because I thereby benefit one person more than another, by giving her a greater chance of receiving some good that she would not have otherwise had. But in the zygotic selection case with twins, Eve's conditional expected well-being depends in no way on our choice; it is higher than Cat's only because we are uncertain about the state of the world, not because (as in a paradigmatically unfair lottery) we gave Eve better chances. Although, in choosing *A*, we cause Cat to have a conditional expected well-being of 1, we do not thereby cause Eve to have a conditional expected well-being of

23. *Ibid.*, 940.

24. Larry S. Temkin, *Inequality* (New York: Oxford University Press, 1993).

65, because she would have had this expectation no matter what we did. We therefore do not benefit Cat and Eve unequally, so we treat no one unfairly. (I say more about the significance of benefiting people in Sec. IV.)

Let me summarize my reasons for doubting that *A* treats Cat unfairly. The mitigatory significance of equal chances given unequal outcomes does not support the claim that unequal chances are independently unfair. It is not reasonable to complain that one was given a worse chance at being well off than others if one had no chance of being worse off than others. And, although causing people to have unequal prospects may count as unfairly benefiting some people more than others, this reasoning does not apply to unaffected individuals, whose expectations of well-being depend in no way on our choice. These considerations count against the hybrid egalitarian verdict in table 5. *A* is not unfair to Cat. So each person's expected well-being, conditional on her existence, is not an ingredient of the currency of egalitarian distribution.

IV. PERSON-AFFECTING PRIORITARIANISM

If I am right that conditional expected well-being is not an ingredient of the currency of egalitarian distribution, or that (even if it is such an ingredient) its marginal contribution to that currency does not decrease, then egalitarians seem unable to account for the intuitive judgment about the risky nonidentity case (table 1). Let me suggest an alternative account.

Consider how Voorhoeve and Fleurbaey motivate their judgment about the case in table 1: "It is more important to provide Ann with the additional benefit of existing at well-being level [49] rather than at [0] if [state 1] is the case than to provide Bob with the additional benefit of existing at [98 + *d*] rather than [49] if [state 2] is the case, because the former would improve Ann's well-being from a lower level."²⁵ This remark is noteworthy for two reasons. First, it seems to appeal to something like a prioritarian concern for increases in well-being at lower levels. Second, this explanation appeals to the greater importance of improving Ann's well-being, or benefiting her. But neither outcome prioritarians nor hybrid egalitarians are particularly concerned with improvements or benefits. To see this, consider a version of the risky nonidentity case in which no one could be benefited, shown in table 6.

In table 6, *A* would not benefit Bob or improve his well-being from a lower level. If we choose *A* and Bob exists, then state 2 obtains, in which case he would not have existed had we chosen *B*. Although we cannot justify *B* to Bob on the grounds that it would maximize his expected well-being, this may not be relevant, because Bob has no complaint against our choice. He is no worse off than he would have been otherwise, and

25. Voorhoeve and Fleurbaey, "Priority or Equality," 931.

TABLE 6
RISKY NONIDENTITY CASE WITHOUT BENEFITS

PROSPECT	STATE 1 (0.5)		STATE 2 (0.5)	
	Ann	Bob	Ann	Bob
<i>A</i>	49			49
<i>B</i>		0	98 + <i>d</i>	

no worse off than anyone else. Since Bob has no complaint against our choice, there is no need to justify it by appeal to anything other than the expected goodness of the prospect from an impartial perspective.²⁶ And, I assume, *B* has greater expected value if $d > 0$, because it offers greater expected (total and average) well-being with no risk of inequality. We should, therefore, choose *B* for any $d > 0$.

Things are different in the original risky nonidentity case (table 1). There, if we choose *B*, Ann has a complaint if state 1 obtains. She can complain that our choice made her worse off than she would have been had we chosen otherwise, in a way that was not in her expected interests. And, crucially, this complaint seems stronger than Bob's, because she is much worse off than Bob would have been had his complaint been satisfied. One's claim to a benefit is stronger, it seems, the worse off one would be were it unsatisfied.

More precisely, according to

Person-Affecting Prioritarianism:

- i) A person has a complaint against our choice if both (*a*) there is some alternative we could have chosen under which she would have fared better and (*b*) our choice fails to maximize her expected well-being, conditional on her existence.
- ii) The strength of a person's complaint against a choice is determined by the extent to which it leaves her worse off than she would have been otherwise, and by how well off she is. A person's complaint is stronger the worse off she is.
- iii) We ought to choose the prospect against which people would have the weakest complaints, discounted by the probability that

26. I assume, following John Broome, *Weighing Lives* (Oxford: Oxford University Press, 2004), 135, that impartial goodness must be assessed anonymously—i.e., in a way that is not sensitive to “the fates of particular persons” (Michael Otsuka, “How It Makes a Moral Difference That One Is Worse Off Than One Could Have Been,” unpublished manuscript). On the view sketched here, benefits to particular people bear on what we ought to do, but not on the goodness of outcomes. See also Jacob Ross, “Rethinking the Person-Affecting Principle,” *Journal of Moral Philosophy* 12 (2015): 428–61.

those complaints arise. If people would have no (or minimally weak) complaints against more than one prospect, then we ought to choose the impartially best of those.

Person-affecting prioritarianism is a *deontic* version of prioritarianism. It does not say which outcomes or prospects are better than others. But I assume, for simplicity, that if an outcome (or prospect) offers greater (expected) total and average well-being than another, then it is better.²⁷

I can now explain how person-affecting prioritarianism handles the cases we have discussed. In the risky nonidentity case with benefits (table 1), person-affecting prioritarianism registers an expected complaint against both prospects. Ann would have a complaint against *B*, were she to exist, because *B* would harm her and fail to maximize her expected well-being conditional on her existence. For the same reason, Bob would have a complaint against *A*, were he to exist. But, for some $d > 0$, the complaint that Ann would have is stronger than the complaint that Bob would have, because Ann would be worse off. We, therefore, have stronger reason to choose *A*. Unlike outcome prioritarianism, however, person-affecting prioritarianism would have us choose *B* if both futures were Ann's, because *B* would then maximize Ann's expected well-being at no one's expense. Ann would then lack a complaint against *B* and would have a complaint against *A*. Person-affecting prioritarianism, therefore, respects the morally significant difference between intrapersonal and interpersonal trade-offs.

In the risky nonidentity case without benefits (table 6), person-affecting prioritarianism registers no complaint against either act. Whatever we choose, no one would be made worse off than she would have been otherwise. Since there is no complaint-based reason to choose either prospect, we should choose the impartially better prospect *B*.

In the zygotic selection cases (tables 3 and 5), person-affecting prioritarianism registers no complaint against *A*, because *A* makes no one worse off than she would have been otherwise. And *A* is better than *B* and *C* because it offers greater expected total and average well-being. Thus, we ought to choose *A*. We should not intervene to bring about a worse life.

27. Person-affecting prioritarians could just as well adopt a more complicated axiology—e.g., to accommodate sensitivity to inequalities in final well-being, or to yield sensible judgments in different-number cases—without affecting my arguments here. Some person-affecting prioritarians might also think, in a pluralistic spirit, that there are other kinds of complaints—e.g., against being made worse off than others, or against being brought into existence with a life not worth living. Some might restrict their conception of complaints to a narrower range of cases in which some agent (e.g., a government) has an independent duty to distribute well-being fairly among some people (e.g., its citizens). And some might allow complaint-based reasons to be outweighed by considerations of goodness in certain cases (following Otsuka, “How It Makes a Moral Difference”).

Person-affecting prioritarianism, therefore, captures the central intuition that motivates hybrid egalitarianism, while avoiding its pitfalls. Of course, person-affecting prioritarianism faces other challenges, but it also has several other virtues; I shall not discuss those other challenges or virtues in detail here.²⁸

Some might wonder whether this view really deserves the name “prioritarianism,” since it does not appeal to the diminishing marginal value of well-being. This seems to me a merely verbal issue, so I won’t say much about it. However, although prioritarianism is typically understood as an impersonal axiological view, on which well-being makes a diminishing marginal contribution to the goodness of outcomes, Parfit himself notes that prioritarianism can be understood either as a view about the goodness of outcomes or as a view about the rightness of acts.²⁹ That most prioritarians opt for the former, axiological characterization can be fully explained by a number of historical factors. For example, many prioritarians have focused exclusively on fixed-population cases, in which one can increase the sum of priority-weighted well-being only by benefiting people. And many philosophers’ interest in prioritarianism derives from a desire to reconcile consequentialism (or restricted consequentialist principles) with a concern for distributive justice.³⁰ It is, therefore, unsurprising that prioritarians have not considered the deontic person-affecting variant proposed here. These historical considerations should not make us doubt that person-affecting prioritarianism captures the fundamental prioritarian insight—that we ought to give priority to the worse-off, not be-

28. One virtue: whereas outcome prioritarianism entails a particularly implausible version of Parfit’s repugnant conclusion (see Campbell Brown, “Prioritarianism for Variable Populations,” *Philosophical Studies* 134 [2007]: 325–61; Matthew D. Adler, “Future Generations: A Prioritarian View,” *George Washington Law Review* 77 [2008]: 1478; Nils Holtug, *Persons, Interests, and Justice* [Oxford: Oxford University Press, 2010]), person-affecting prioritarianism has no such implication, because one cannot have a complaint if one never exists. One challenge: person-affecting prioritarianism registers no reason to improve the final well-being of the worse-off when doing so fails to increase anyone’s expected well-being (see Voorhoeve and Fleurbaey, “Priority or Equality,” 936). Some person-affecting prioritarians, however, might give weight to inequality in various ways, as mentioned in n. 27. Others might find it wrong to sacrifice the expected well-being of each person for the sake of some impersonal good, such as equality. I am inclined to favor the latter, less concessive route.

29. Derek Parfit, “Equality and Priority,” *Ratio* 10 (1997): 202–21, 213 n. 18. Another deontic version of prioritarianism is proposed by Andrew Williams, “The Priority View Bites the Dust?,” *Utilitas* 24 (2012): 315–31. But, on Williams’s view (and not on mine), the strength of our reasons to benefit a person can depend on whether and how our act would affect other people. And Williams does not ground his deontic view in an appeal to person-affecting considerations.

30. See, e.g., Samuel Scheffler, *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*, rev. ed. (Oxford: Clarendon, 1994), chap. 2; and Brad Hooker, *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Clarendon, 2003), chap. 2.

cause they are worse off than others, but rather because (in Parfit's words, again) "we have stronger reasons to benefit people the worse off these people are." As I see it, outcome prioritarianism is just an implausible overgeneralization of this insight. It regards our reasons to benefit the worse-off as merely a special case of our more general reasons to increase the sum of priority-weighted well-being. But the sum of priority-weighted well-being is not, in itself, a proper object of moral concern. Priority is owed to people and to their claims, not to quantities of well-being.

V. CONCLUSION

Unlike outcome prioritarianism, Voorhoeve and Fleurbaey's hybrid egalitarianism respects the morally significant difference between intrapersonal and interpersonal trade-offs. But, as I showed in Section II, hybrid egalitarianism violates statewise dominance, the weak principle of personal good, or the tight connection between well-being and personal good. And, as we saw in Section III, it implausibly recommends prospects that are certain to increase inequality in final well-being, in order to prevent unobjectionable inequalities in expected well-being. Person-affecting prioritarianism avoids these problems while also respecting the morally significant difference between intrapersonal and interpersonal trade-offs.