

Fair Pay

Kristi A. Olson

Introduction

Consider the following inequality: Helen and Ivan are from the same social background with identical talents and resources. They initially do the same work for the same pay. Yet, when faced with the identical opportunity to work overtime in exchange for additional income, Helen accepts while Ivan declines. As a result, at the end of the year, Helen has more money than Ivan. Nonetheless, the inequality between them strikes most people as fair.

If you are not an egalitarian, you probably will not find this scenario puzzling at all. But if you are an egalitarian and your default assumption is that resources should be equal, the fairness of the inequality requires an explanation.¹ Even to egalitarians the inequality between Helen and Ivan seems fair, *but why is it fair?*

The conventional answer appeals to choice. In order to explain why some inequalities but not others are fair, choice-sensitive egalitarians (also known, more pejoratively, as luck egalitarians) introduced the following distinction: Inequalities that are not attributable to choice are unfair and, as such, require rectification. Thus, we should rectify inequalities attributable to differences in natural talents. In contrast, inequalities that are attributable to choice are fair and, as such, do not require rectification. The reason we should not take from Helen to redistribute to Ivan, then, is because the inequality between them was attributable to their freely-made choices.

What is puzzling about this answer, however, is that it seems to ignore an equally, if not more, plausible solution. After all, there seem to be (at least) two possible explanations for why the inequality between Helen and Ivan is fair. According to the first explanation, the inequality is fair because it is attributable to choice: Helen *chose* to work overtime; Ivan did not. In contrast, according to the second explanation, the inequality is fair because of the differences in their work: Helen *worked overtime*; Ivan did not. The additional income, according to the second explanation, merely compensates Helen for her additional burdens and thus there is no inequality, all things considered.

¹ As this sentence reveals, my target audience in this essay is resource egalitarians. If you believe that individuals' welfare levels—as opposed to their resources—should be equal, then the example is underspecified: we are told nothing about Helen's and Ivan's welfare levels. Nonetheless, although the paper is primarily directed at resource egalitarians, much of what I say applies to egalitarians more broadly.

Both explanations seem capable of justifying the inequality between Helen and Ivan. Yet, since the two justifications can pull in opposite directions, egalitarians cannot simply appeal to both. The task of this essay, then, is to adjudicate between them. In Part I of this essay, I show that choice-sensitive egalitarians have long conflated the two justifications and, as a result, neither justification has been judged on its own merits. In Part II, I distinguish two interpretations of the choice justification and explain why neither interpretation is able to justify disparities in pay. Finally, in Part III, I flesh out an alternative justification for inequalities in pay based on differences in work.

Before turning to my argument, however, let me briefly explain why egalitarians need to be able to justify inequalities in pay in the first place. The judgment that the inequality between Helen and Ivan is fair enjoys nearly universal acceptance across the political spectrum. Indeed, the anti-egalitarian right use precisely such examples to motivate their cause.² Thus, if egalitarianism is to have any political plausibility at all, egalitarians cannot insist on equal resources without exception. Rather, they must present a more nuanced view: You don't have to abandon your egalitarian ideals just because you think it is fair for Helen to have more. It would not be enough, however, for egalitarians simply to stipulate that the particular inequality between Helen and Ivan is fair. After all, unless they give us a principled explanation, egalitarianism is incapable of providing guidance. How can we judge one distributive scheme to be more or less egalitarian than another if we do not know which inequalities egalitarianism condemns and which it embraces? Thus, if egalitarianism is to have any plausibility, egalitarians owe us a principled explanation. The purpose of this essay, then, is to provide that explanation. I begin by showing why the conventional answer fails.

Part I.

A. Although the foundations of choice-sensitive egalitarianism are sometimes traced back to John Rawls,³ the distinction between inequalities that arise from an individual's choices and inequalities that arise from an individual's place in the distribution of natural talents was first made prominent by Ronald Dworkin.⁴ The fundamental idea Dworkin introduced was hailed as a great

² See, e.g., <http://clevelandteapartypatriots.blogspot.com/2010/01/ant-grasshopper.html>; <http://www.warroom.com/antgh.php>.

³ See, e.g., Kymlicka, *Contemporary Political Philosophy*, pp. 53-75. For arguments denying that the foundations of choice-sensitive egalitarianism can be traced back to Rawls, see Samuel Scheffler, "What is Egalitarianism?" *Philosophy & Public Affairs* 31 (2003): 5-39 and Samuel Freeman, *Justice and the Social Contract* (Oxford: Oxford Univ. Press 2007), pp. 111-142.

⁴ Dworkin, "What Is Equality? Part 2. Equality of Resources," *Philosophy & Public Affairs* 10 (1981): 283-345, reprinted in Dworkin, *Sovereign Virtue* (Cambridge, Mass: Harvard Univ. Press, 2000), pp. 65-119.

step forward. In the oft-quoted words of G.A. Cohen, “Dworkin has, in effect, performed for egalitarianism the considerable service of incorporating within it the most powerful idea in the arsenal of the anti-egalitarian right: the idea of choice and responsibility.”⁵ Yet, as I will argue below, this praise is misplaced. To the contrary, much of the appeal of Dworkin’s account stems from the fact that he conflated choice and hard work. In the remainder of Part I, I explain how the conflation arises in Dworkin’s account and then show that the same conflation pervades choice-sensitive egalitarianism. In order to explain how the conflation arises in Dworkin’s account, however, I first need to lay out his account. As we shall see, this is all the more necessary since his account has been widely misunderstood.

Dworkin set out to design a redistributive scheme that would “neutralize the effects of differential talents”, yet “preserve the consequences of [individual choice].”⁶ In order to do this—or at least to achieve the closest approximation—Dworkin proposed a tax and redistribution scheme based on the following thought experiment.⁷ Suppose we know what talents we possess but are ignorant of their economic value. Before learning their economic value, we are given the opportunity to purchase underemployment insurance as follows: If our maximum earning capacity (as determined by the market) falls below the coverage level, the insurance pays us the difference between our maximum earning capacity and the coverage level, minus the premium we must pay for the insurance.⁸ On the other hand, if our maximum earning capacity exceeds the coverage level, we must still pay the premium, but the insurance pays nothing. Although we can purchase whatever level of coverage we choose, the higher the level of coverage, the higher the premium.

Since it is not possible for (adult) individuals to choose underemployment insurance before knowing the economic value of their talents, the insurance decision is intended to be merely hypothetical. What we need to know, then, is what level of insurance individuals would *hypothetically* choose. The following example illustrates how the decision might be made: Suppose that prior to knowing the economic value of her talents, Anna is trying to decide whether to purchase underemployment insurance and, if so, at what level of coverage. She reasons as follows: Suppose my earning capacity turns out to be quite low, say \$10,000. There are many things I would not be able to do with such a low maximum earning capacity. Indeed, I might not be able to support

⁵ Cohen, “On the Currency of Egalitarian Justice,” *Ethics* 99 (1989): 906-44, at p. 933.

⁶ Dworkin, *Sovereign Virtue*, p. 91.

⁷ *Ibid.*, pp. 65-119.

⁸ Importantly, the insurance policy pays only the difference between the coverage level and the individual’s maximum earning capacity, not the difference between the coverage level and what the individual actually earns. Thus, individuals are held responsible for their choice to earn less than their maximum earning capacity. See, e.g., Dworkin, *Sovereign Virtue*, p. 94 (“Insurance is provided against failing to have an opportunity to earn whatever level of income, within the projected structure, the policyholder names, in which case the insurance company will pay the policyholder the difference between that coverage level and the income he does in fact have an opportunity to earn.”)

myself. Thus, Anna concludes that she should purchase insurance. But what level of coverage should she purchase? Suppose Anna narrows down the possibilities to two options: insuring at the level of \$200,000, with an \$80,000 premium, or insuring at the level of \$60,000, with a \$5,000 premium. In trying to decide between them, Anna imagines the following scenario: Suppose I insure at the level of \$200,000 and then find out I could be a corporate lawyer for \$200,000 or a philosophy professor for \$50,000. Even though I would hate to be a corporate lawyer, I would be forced to work in that capacity at least some of the time just to pay the premium. On the other hand, if I insure at \$60,000 and then turn out to have those same earning capacities, I would still be able to pursue a career as a philosophy professor. Anna then reasons as follows: Insuring at the \$60,000 level ensures an adequate minimum if my earning capacity turns out to be very low. At the same time, however, because the premium is relatively low, there is very little chance that I would have to pursue my maximum earning capacity just to pay the premium.⁹ Thus, Anna chooses to insure at the level of \$60,000.

Dworkin uses a similar line of reasoning to argue that, like Anna, individuals would choose to insure, but only at relatively low levels of coverage. Dworkin then proposes a tax and redistribution scheme based on this level of insurance.¹⁰ Under Dworkin's tax and redistribution scheme, individuals who are worse off because their maximum earning capacities are below the coverage level (and thus not the result of choice) receive compensation. In contrast, individuals who are worse off because they choose to work below their maximum earning capacity must bear the cost. To illustrate: Suppose we decide that individuals would purchase insurance at the level of

⁹ Anna's reasoning is actually not very good. What she should care about is not whether she is forced into her maximum earning capacity just to pay the premium. Rather, what she should care about are the choices she is left with under her various insurance options. This, however, depends on some empirical facts that Anna does not consider. To illustrate the problem, suppose Anna turns out to have two earning capacities: she can earn \$60,000 as a professor or \$25,000 as a novelist. If Anna insured at the level of \$60,000 with a \$5,000 premium and if the \$20,000 remaining to her as a novelist after she pays the insurance premium is not enough to meet Anna's needs, then Anna would be forced to work as a professor. In contrast, if she had insured at the level of \$200,000, she would be free to be a novelist since, in addition to her \$25,000, she would receive \$60,000 (the difference between \$200,000 and \$60,000 minus the \$80,000 premium). As a result, Anna cannot simply conclude that she should not insure at the higher level because doing so might force her to work at her maximum earning capacity. After all, she might be forced to work in her maximum earning capacity as a result of insuring at the lower level. Instead, Anna first needs to ascertain empirical facts about the likely distribution of earning potentials. On this issue, see Dworkin, *Sovereign Virtue*, pp. 98-9. For our purposes, however, we can set this issue aside.

¹⁰ The way in which Dworkin's insurance scheme is supposed to be sensitive to individuals' choices is, in fact, ambiguous. On the one hand, it might be the individual's (hypothetical) choice to purchase insurance that makes the resulting distribution choice-sensitive. On the other hand, however, the insurance scheme is itself designed to produce an outcome that reduces the effects of differential talents, while preserving the consequences of occupational choice. Thus, choice arguably comes into play twice in Dworkin's insurance solution: the (hypothetical) choice of insurance coverage and the (actual) choice of occupation. On this issue, see Robert van der Veen, "Equality of Resources: Procedures or Outcomes?" *Ethics* 113 (2002): 55-81.

\$60,000 with a \$5,000 premium. If my maximum earning capacity turns out to be \$40,000, I will receive a payout of \$15,000 (the difference between the coverage level and my maximum earning capacity, minus the premium) no matter what I choose to do. As a result, if I choose to work at my maximum earning capacity, my adjusted earnings will be \$55,000 (\$40,000 plus \$15,000). On the other hand, if I choose to work at half my maximum earning capacity, my adjusted earnings will be \$35,000 (\$20,000 plus \$15,000). Thus, even though I receive the same insurance payout regardless of what I choose to do, my total income still depends on my choices.

B. Dworkin's insurance scheme has been widely interpreted as instantiating the intuitively appealing idea that it is not unfair for those who choose to work harder to receive more.¹¹ Yet, this is not what Dworkin's insurance scheme actually tracks. The choice to earn more relative to one's earning capacity is different from the choice to work harder and Dworkin's insurance scheme tracks the former and not the latter.

To illustrate the difference: Suppose Brigitte and Clara each have two options. Brigitte could earn \$30,000 as a nurse or \$60,000 as a plumber. Clara could earn \$30,000 either as a novelist or as a nurse. Both choose to be nurses because they find the work more meaningful than their other options. Brigitte, however, chooses to work full-time, while Clara chooses to work part-time and spends the remainder of her time playing video games. They work on the same floor of the same hospital doing the same work, except that Brigitte works five shifts per week (earning \$30,000), while Clara works only three (earning \$18,000).

The choice to work harder, then, is different from the choice to earn more relative to one's earning capacity. Although Brigitte chooses to work harder than Clara, Clara chooses to earn more relative to her earning capacity than Brigitte.¹² Since they pull in opposite directions—the choice to work harder favors Brigitte, the choice of earnings relative to capacity favors Clara—Dworkin's insurance scheme cannot track both. Yet, instead of tracking the choice to work hard, as his insurance scheme is generally interpreted, Dworkin's redistributive scheme in fact tracks the choice to earn more relative to one's earning capacity. If Brigitte and Clara are assumed to have purchased \$60,000 insurance with a \$5,000 premium, Clara will receive \$25,000 from the

¹¹ See, e.g., Kwame Anthony Appiah, "Equality of What?" *The New York Review of Books*, Apr. 26, 2001 (summarizing Dworkin as saying "you can only have the riches if you accept the drudgery."); Eric Rakowski, *Equal Justice*, p. 99 (interpreting Dworkin as endorsing "the fundamental principle ... that people ought to start life with an equal bundle of resources and that equality should be maintained over time except insofar as somebody works less hard or accepts more desirable work ...").

¹² Specifically, Clara chooses to earn 60 percent of what she is capable of earning, while Brigitte chooses to earn only 50 percent of what she is capable of earning.

insurance policy (the difference between the coverage level of \$60,000 and her maximum earning capacity of \$30,000 minus the \$5,000 premium). Brigitte, on the other hand, receives nothing (the difference between the coverage level of \$60,000 and her maximum earning capacity of \$60,000) but must still pay the \$5,000 premium. Thus, as Table 1 illustrates, under Dworkin's hypothetical insurance scheme, Brigitte ends up with significantly less income despite the fact that she works considerably longer. Specifically, Clara's adjusted income is \$43,000, while Brigitte's is only \$25,000.¹³

Table 1. Brigitte's and Clara's Earnings

	Unadjusted Earnings	Insurance Adjustment (\$60K Insurance, \$5K Premium)	Adjusted Earnings
Brigitte (full-time)	\$30,000	-\$5,000	\$25,000
Clara (60% time)	\$18,000	+\$25,000	\$43,000

Thus, although Dworkin seems to assume that his insurance scheme rewards the individual who chooses to work harder,¹⁴ this is not what his insurance scheme actually tracks. Instead, Dworkin's insurance scheme tracks how much the individual chooses to earn relative to what she is capable of earning. Examples in which the individual who chooses to work harder receives more, then, are misleading: They might—and, indeed, did—cause readers to think the insurance scheme tracked hard work.¹⁵

C. Dworkin might have been the first choice-sensitive egalitarian to conflate choice and hard work, but he was certainly not the last. To the contrary, the reason the conflation of choice and hard work has been so easy to overlook is not because it is rare, but rather because it is exceedingly common. To illustrate the extent to which the conflation pervades choice-sensitive egalitarianism, consider the following examples: Will Kymlicka motivates choice-sensitive egalitarianism with an

¹³ There are in fact multiple distributions that track individuals' choice of earnings relative to their earning capacity. For example, we could allocate Clara sixty percent of \$60,000, since she only works at sixty percent capacity, leaving her with \$31,000 after paying the premium instead of \$43,000. For my purposes, however, the details do not matter. What matters for my purposes is simply that Dworkin's account sometimes justifies inequalities in which the harder-working individual receives less. Thus, his account does not track the choice to work harder.

¹⁴ In motivating the need for his hypothetical insurance solution, Dworkin sometimes refers to the individual's choice to work in more rather than less profitable ways (relative to her capacity), but he also refers to the individual's choice to work hard or to undertake work no one else wants to do. For example, in one sentence, he invokes the individual's capacity to earn more: "Suppose that two people have very different bank accounts, in the middle of their careers, because one decided not to work, or not to work at the most lucrative job he could have found, while the other single-mindedly worked for gain." Yet, in the very next sentence, he appeals to hard work: "Or because one was willing to assume especially demanding or responsible work, for example, which the other declined." Dworkin, *A Matter of Principle*, p. 206. For other references to hard work, see Dworkin, *Sovereign Virtue*, p. 89; Dworkin "Equality, Luck and Hierarchy," *Philosophy & Public Affairs* 31 (2003): 190-98, at p. 193.

¹⁵ See n. 11.

example in which one individual (the gardener) chooses to work “longer hours” while the other individual (the tennis player) chooses to pursue “more leisure.”¹⁶ Michael Otsuka does essentially the same: suppose “one person freely chooses to spend his days surfing and living off whatever food he can gather at his leisure ..., whereas the other freely chooses to spend his days cutting down his trees in order to construct a magnificent house and garden”¹⁷ Stuart White uses an example involving longer hours of work: “if Smith has more income than Jones because he chooses to work longer hours then [Dworkin’s] principle of ambition-sensitivity requires that we respect this inequality.”¹⁸ Larry Temkin similarly writes that it is not objectionable “when one person is worse off than another due to her own responsible choices, say to pursue a life of leisure”¹⁹ Miriam Cohen Christofidis explicitly invokes hard work: “True, in virtue of working hard, that person might be able to afford to have nice things, things that I would like to have, but, *ex hypothesi*, I too was equally able to work hard, and thereby acquire those things, so I cannot reasonably object to this inequality between us.”²⁰ G.A. Cohen, in explaining what motivated egalitarians to appeal to choice in the first place, suggests that they were responding to objections to strict equality invoked by images of the hardworking ant and the idle grasshopper: “why should those, like the grasshopper and the ant, with exactly the same initial advantages, and who merely *chose* differently, be forced back to equality if an inequality ensues?”²¹ Serena Olsaretti similarly explains: “By qualifying the demands of just equality ... luck egalitarianism is able to escape familiar anti-egalitarianism criticisms like the claim that it panders to the undeserving poor at the expense of hard-working citizens.”²² Falling into the same trap, Kok-Chor Tan distinguishes inequalities that arise as a result of the social class one is born into from inequalities that arise as a result of ambition and “hard

¹⁶ Kymlicka, *Contemporary Political Philosophy* (Oxford: Oxford Univ. Press, 2002), pp. 72-73. See also Kymlicka, “Left-Liberalism Revisited,” in *The Egalitarian Conscience: Essays in Honour of G.A. Cohen*, ed. Christine Sypnowich (Oxford: Oxford Univ. Press, 2006), pp. 9-35, at p. 18 (describing choice-sensitive egalitarianism in terms of the choice to “give up leisure” or “to work fewer hours”).

¹⁷ Otsuka, “Equality, Ambition and Insurance,” *Aristotelian Society Supplementary Volume 78* (2004): 151-166, at p. 152. Otsuka’s example introduces an additional possible justification since we might think that an individual has a right to use without interference the products she creates. For our purposes, however, we can set this additional complexity aside.

¹⁸ White, *Equality* (Cambridge: Polity Press, 2007), p. 80.

¹⁹ Temkin, “Justice, Equality, Fairness, Desert, Rights, Free Will, Responsibility, and Luck,” in Carl Knight and Zofia Stemplowska (eds.), *Responsibility and Distributive Justice* (Oxford: Oxford Univ. Press, 2011): 51-76, at p. 62.

²⁰ Christofidis, “Talent, Slavery and Envy in Dworkin’s Equality of Resources,” p. 269.

²¹ Cohen, “Luck and Equality: A Reply to Hurley,” *Philosophy and Phenomenological Research* LXXII (2006): 439-446, at 444 (his emphasis). Cohen’s explanation for the motivation behind choice-sensitive egalitarianism is, admittedly, quite puzzling given Cohen’s view, expressed elsewhere, that there is no inequality if individuals receive more simply to compensate for additional burdens. Yet, the very fact that Cohen sees the example of the ant and grasshopper as the motivation for choice-sensitive egalitarianism suggests how deeply ingrained the conflation really is.

²² Olsaretti, “Responsibility and the Consequences of Choice,” *Proceedings of the Aristotelian Society* CIX (2009): 165-88, at 166.

work.”²³ John Roemer, although more subtle, uses “effort” to denote both the individual’s choice to work in more rather than less profitable ways and how hard the individual works.²⁴ Alexander Cappelen and Bertil Tungodden appear to do the same.²⁵

Indeed, even the critics of choice-sensitive egalitarianism have accepted and reinforced the characterization of choice-sensitive egalitarianism as rewarding hard work. Jonathan Wolff, for example, defers to Kymlicka’s example of the hard working gardener and the leisure pursuing tennis player.²⁶ Colin MacLeod motivates choice-sensitive egalitarianism by appealing to differential labor: “Suppose you choose to brew some beer while I choose to spend the days watching TV.”²⁷ Samuel Scheffler refers to longer hours of work: “If some people make more money than others because they choose to work longer hours”²⁸ And the list continues.

Of course, it is not a mere coincidence that the stock examples used to motivate choice-sensitive egalitarianism all involve an inequality favoring the individual who works harder. These examples are used precisely because they present the strongest illustration of justified inequalities. Yet, the very fact that these examples present the strongest illustration of justified inequalities undermines the argument that *choice*—as opposed to hard work—justifies the inequality. Nonetheless, my argument in Part I is merely that choice and hard work have been conflated. For all I have said thus far, it is still *possible* that choice justifies inequalities. I turn to that argument in Part II.

Part II.

A.

Before we can examine the claim that choice justifies inequalities, we first need to distinguish two quite different interpretations of the way in which choice justifies inequalities. According to what I will call the capacity interpretation, what justifies an inequality is an

²³ Tan, “A Defense of Luck Egalitarianism,” *The Journal of Philosophy* CV (2008): 665-690, at p. 669. See also p. 684 (equivocating about the meaning of “effort”).

²⁴ Roemer, “Defending Equality of Opportunity,” p. 261 and 279 (defining effort as “autonomously chosen action—within the individual’s control—which, if expended in greater amounts, will increase the degree to which the individual acquires the objective” but also equating the individual’s autonomously chosen action with “exerting great effort” and “work[ing] hard”). Susan Hurley also observed that there is an equivocation in Roemer’s account, but she attributes the equivocation to desert (where desert is not simply tracking the burdensomeness of the work). Hurley, *Justice, Luck, and Knowledge* (Cambridge, Mass: Harvard Univ. Press, 2003), p. 182.

²⁵ Cappelen and Tungodden, “Rewarding Effort,” *Economic Theory* 31 (2009): 425-441 at p. 426.

²⁶ Wolff, “Fairness, Respect, and the Egalitarian Ethos,” *Philosophy & Public Affairs* 27 (1998): 97-122, at pp. 99-100.

²⁷ MacLeod, *Liberalism, Justice, and Markets*, p. 10.

²⁸ Scheffler, “What is Egalitarianism?” p. 6. Scheffler’s example is especially revealing since he uses the example to illustrate the limited extent to which choice-sensitive egalitarianism happens to agree with the prevailing morality.

individual's choice to earn less than she is capable of earning relative to the choices of others.²⁹ This is the justification we saw at work in Dworkin's hypothetical insurance scheme.³⁰ The reason Brigitte's adjusted earnings were less than Clara's was because Brigitte chose to earn less, relative to her earning capacity, than Clara did.

In contrast, according to what I will call the estoppel interpretation, what justifies an inequality is the fact that the individual could have made the identical choice as the better off individual, but did not, and thus cannot legitimately complain. To illustrate, suppose you choose to be a corporate lawyer for \$200,000 while I choose to be a philosophy professor for \$50,000. If I also could have been a corporate lawyer for \$200,000, then, according to the estoppel interpretation, I cannot legitimately complain. After all, I too could have made that same exact choice. On the other hand, if you choose to be a professional basketball player with a million dollar income and I could not have made that same exact choice, then the inequality between us is not justified. The estoppel interpretation, then, is more restrictive than the capacity interpretation: it only justifies inequalities when the worse off individual could have chosen the same option (or, perhaps, an option at least as good³¹) as the option chosen by the better off individual.

Although the two interpretations are rarely distinguished, since the two interpretations justify different inequalities, choice-sensitive egalitarians cannot simply appeal to both. Each of the interpretations, then, must be considered on its own merits. Thus, I address each interpretation in isolation below.

B. To begin my critique of the capacity justification, let us first consider an example that Dworkin himself raises.³² In that example, Deborah and Ernest choose to do the same job but Deborah, unlike Ernest, could earn much more as a movie star. Because Ernest's maximum earning capacity is below the coverage level, Ernest receives compensation from the insurance policy.

²⁹ The capacity interpretation can also be put in terms of contribution: an inequality is justified if the individual chooses to contribute less than she is capable of contributing, relative to the choices of others (where contribution is measured by market value).

³⁰ The capacity interpretation is also at work in the distributive schemes of John Roemer, Stuart White, and Miriam Cohen Christofidis. See, e.g., John Roemer, "Defending Equality of Opportunity," *The Monist* 86 (2003): 261-82; Roemer, et al., "To What Extent Do Fiscal Regimes Equalize Opportunities for Income Acquisition among Citizens?" *Journal of Public Economics* 87 (2003): 539-65; Roemer, *Egalitarian Perspectives: Essays in Philosophical Economics* (Cambridge: Cambridge Univ. Press, 1994), pp. 179-96; Stuart White, "The Egalitarian Earnings Subsidy Scheme," *British Journal of Political Science* 29 (1999): 601-622; and Miriam Cohen Christofidis, "Talent, Slavery and Envy in Dworkin's Equality of Resources," *Utilitas* 16 (2004): 267-87.

³¹ This of course presupposes that we have the capacity to determine when an option is at least as good. I discuss how such judgments can be made in Part III.

³² Dworkin, *Sovereign Virtue*, p. 99. Dworkin attributes the example to Thomas Scanlon.

Deborah, in contrast, must pay the premium, but receives nothing. As a result, Deborah ends up with less pay than Ernest, despite doing identical work.

By virtue of holding the work they do constant, the example isolates choice and, as a result, is the perfect test case for the claim that choice (specifically, the choice of earnings relative to one's capacity) and not hard work exonerates inequalities. Yet, instead of embracing the example and putting it forward as motivation for his view, Dworkin seems slightly embarrassed by it. He begins by describing the example as an "apparent anomal[y]".³³ In order to explain away its apparent anomalousness, Dworkin first argues that, if Deborah and Ernest had insured at low levels of coverage (as Dworkin argues people will), the difference between them would be a modest one.³⁴ He then claims that, if the difference were a modest one, the difference "would (arguably) then be an appropriate mark of the fact that Deborah had an option Ernest did not."³⁵

The problem with Dworkin's response, however, is that, in attempting to explain away the apparent anomalousness, he illicitly assumes the circumstances most conducive to his view. Specifically, Dworkin assumes that the income inequality between Deborah and Ernest is a modest one and his specification of the example at least suggests that Deborah's higher-paying option is a good one and that Deborah has more and better options than Ernest. Yet, Dworkin is not entitled to any of these assumptions. To illustrate, imagine that Deborah is capable of a job that pays \$60,000, but the work is extremely dangerous and oppressive. Her only other option pays only \$30,000 but, let us stipulate, it is still the all things considered better option. Ernest, in contrast, is incapable of the onerous job that pays \$60,000, but he has many highly enjoyable options all of which pay \$30,000, including the one available to Deborah.

Despite the fact that Ernest has more and better options than Deborah, the capacity interpretation would sanction an inequality favoring Ernest. Moreover, the inequality between them could be considerable. As Table 2 illustrates, even at if Deborah and Ernest are assumed to have insured at relatively low levels (say, \$60,000 with a \$5,000 premium), if they choose to do the same work, Ernest will receive \$55,000 while Deborah receives \$25,000.

Table 2. Deborah's and Ernest's Adjusted Earnings

	Unadjusted Earnings	Insurance Adjustment (\$60K Insurance, \$5K Premium)	Adjusted Earnings
Deborah (onerous job)	\$60,000	-\$5,000	\$55,000
Deborah (easy job)	\$30,000	-\$5,000	\$25,000
Ernest (easy job)	\$30,000	+\$25,000	\$55,000

³³ Ibid.

³⁴ Ibid.

³⁵ Ibid.

Deborah, then, is faced with the following choice: If she chooses to work in her maximum earning capacity, her income will be the same as Ernest's, but her work will be significantly worse. On the other hand, if she chooses to do the same work as Ernest, the work will be the same but her income will be significantly less. Thus, no matter what she chooses, Deborah will be significantly worse off than Ernest.³⁶

Yet, as bad as those options might seem, matters could be even worse. To illustrate: Suppose that Deborah has only the following two options available to her: an extremely onerous job that pays \$60,000 or a slightly less onerous job that pays \$30,000. Ernest, in contrast, has a great many highly-enjoyable jobs that all pay \$30,000. Under these circumstances, Deborah faces an even worse choice: If she chooses to work in her maximum earning capacity, her income will be the same as Ernest's, but her work will be significantly worse. On the other hand, however, if she chooses the slightly less onerous job, her work will still be worse than Ernest's, but, in addition, her income will be significantly less.

Presumably few egalitarians would deem the inequality between Deborah and Ernest to be fair.³⁷ Indeed, even with his illicit assumptions, Dworkin still seemed uncomfortable with the inequality. After attempting to explain away the apparent anomalousness, his final response to the example was to dismiss it: "In any case, this unfairness, if it is unfairness, would disappear in any plausible translation of the hypothetical insurance market into an actual tax scheme"³⁸ If the example had no special significance, such a dismissal might be plausible. The problem, however, is that the example *does* have special significance: In the only example Dworkin addresses in which he holds labor burdens constant—and thus the only example that could be used to motivate his claim that choice and not hard work justifies inequalities—he himself acknowledges that the inequality might be unfair.³⁹ And the reason the inequality is unfair is precisely because it ignores how hard the work is.

³⁶ Dworkin might object that especially dangerous jobs must be excluded from the application of his insurance scheme. That, however, does not solve the problem. At most, it narrows the range of cases to which the problem applies. (Perhaps Deborah's job is much more onerous but only a little more dangerous.) Moreover, the very need for such an exception illustrates the problem with the capacity interpretation.

³⁷ Egalitarians might be more likely to deem the inequality between Deborah and Ernest fair if Deborah and Ernest had *actually* (as opposed to hypothetically) selected that insurance policy. That, however, cannot salvage the capacity interpretation. First, for the very reasons Dworkin notes, it is impossible for individuals actually to choose underemployment insurance before knowing the economic value of their talents. Second, even if the choice of underemployment insurance were to justify the inequality, that justification is entirely distinct from the capacity interpretation.

³⁸ Dworkin, *Sovereign Virtue*, p. 99.

³⁹ A more promising example for choice-sensitive egalitarians is the inequality between the voluntarily and involuntarily unemployed. Although they do the same work—i.e., none—many people think the voluntarily unemployed should receive less than the involuntarily unemployed. Nonetheless, there are several reasons why choice-sensitive egalitarians should not rest their case on this example. First, in order to show that choice is doing the justificatory work, examples comparing

The problem with the capacity interpretation, then, is not merely that it sanctions inequalities that few egalitarians would endorse. The more serious problem is that these inequalities—precisely because they do not track hard work—are the paradigm examples of the view. Thus, the very feature that makes the example of Deborah and Ernest capable of showing that choice—as opposed to hard work—justifies inequalities also makes the inequality between them implausible. And that is a serious problem for proponents of the capacity interpretation.

C. The estoppel interpretation, at least at first glance, seems more promising. It avoids the particular problem associated with the capacity interpretation. According to the estoppel interpretation, the inequality between Deborah and Ernest would not be justified because, unlike Ernest, Deborah did not have the option of receiving \$55,000 while pursuing a relatively enjoyable job. Thus, the estoppel interpretation would not sanction the inequality between them.

The estoppel interpretation, however, runs into a different problem. The problem begins from the fact that there are two different roles that a justification might play. The first concerns the *bundling* of benefits and burdens. The second concerns the *assignment* of bundles. The problem for the estoppel interpretation is that choice (specifically, choice between identical option sets) justifies the assignment of bundles. Yet, it is the justification for the bundles themselves that justifies the disparities in pay.

The following example illustrates:⁴⁰ Suppose a divorcing couple, unable to reach an agreement, request that the judge divide what property remains after their debts have been paid. The wife has such a strong sentimental attachment to the hockey equipment she used when she

the voluntarily and involuntarily unemployed must hold everything else—including the individuals' leisure time activities and total resources—constant. Thus, if we want to show that choice is doing the justificatory work, it would be illegitimate to compare an impoverished individual too sick to work to a wealthy and healthy retired CEO who spends his time surfing. After all, these differences might be doing some or all of the justificatory work. A second reason choice-sensitive egalitarians should not rely on this example is that, although it might be fair for someone with fewer (or worse) options to receive more income (in compensation for their lack of options), that is distinct from the claim that choice justifies inequalities. Since the involuntarily unemployed individual by definition has fewer options than a voluntarily employed individual, this too might be doing some of the justificatory work and not choice. To minimize the possible distorting effect of this difference, we should imagine a voluntarily unemployed individual with very few employment options and whose options are undesirable. The third worry is that we might think individuals who are able to work have some obligation to do so. This could explain why we think the voluntarily unemployed should receive less (or nothing at all). That, however, is also distinction from the claim that choice justifies inequalities. Thus, once we control for all these other variables, my suspicion is that the example is perceived unfairness is likely to disappear. In any case, however, choice-sensitive egalitarians would also need to explain why the Deborah/Ernest example is not equally persuasive. If choice is really doing the justificatory work, why aren't both examples compelling?

⁴⁰ For similar examples, see my "Our Choices, Our Wage Gap?" *Philosophical Topics* 40 (2012): 45- 61, at pp. 48-49 and Joseph Heath, "Dworkin's Auction," *Politics, Philosophy & Economics* 3 (2004): 313-335, at pp. 316-17. Although Heath was making a different point, his example is similar to the one I use here.

taught her now-grown daughters to play that she would prefer a bundle with the hockey equipment to any other bundle. The husband in contrast has an equally strong sentimental attachment to their ping pong table. The judge, knowing their preferences and favoring the husband, gives the wife first choice of the following bundles: {the hockey equipment and a few worthless trinkets} or {the ping pong table and the lion's share of valuable property}. Not knowing her husband's preferences and thus unable to engage in arbitrage, the wife chooses the bundle with the hockey equipment; the husband is delighted with the result. When the wife suggests that the division was unfair, the judge explains to her that, according to the estoppel interpretation, since she had the same options as her husband, she cannot legitimately complain.

The judge's claim that she cannot legitimately complain, however, is only partially true. In particular, it is true that she cannot legitimately complain that she has her bundle and not his since, by hypothesis, she could have chosen his bundle. That, however, is not her complaint; her complaint is not about the *assignment* of bundles. Rather, her complaint is that the bundles were themselves unfair.⁴¹ The estoppel interpretation, then, might be able to justify the *assignment* of bundles, but it does not justify the bundles themselves. Yet, when it comes to justifying disparities in pay, what we need is a justification for the bundles themselves.

We can now see why the examples in which the individual who chooses to work harder receives more present the strongest illustration of justified inequalities. By appealing to both choice and hard work, choice-sensitive egalitarians simultaneously justify the bundling of benefits and burdens and the assignment of bundles. Yet, by appealing to examples that conflate the two justifications, egalitarians obscured the distinct role that each justification plays. In doing so, they also obscured the fact that, although choice under the estoppel interpretation can justify the assignment of bundles, it does not—and indeed cannot—establish that the bundles themselves are fair.

The problem, then, is not merely that choice by itself is insufficient to justify inequalities, as some choice-sensitive egalitarians have recognized.⁴² Rather, the problem is that choice, under the estoppel interpretation, cannot tell us whether a disparity in pay is fair or unfair. The different roles—assignment versus bundling—explains why this is so. Choice only speaks to the assignment of bundles. If we want to know whether a disparity in pay is fair or unfair, we need a justification for the bundles themselves. I take up this task in Part III.

⁴¹ The wife might also have a legitimate complaint against the judge for intentionally discriminating against her. The unfairness, however, does not depend on discrimination. Suppose, for example, that the judge simply made a mistake in filling out the form. The wife would still have a legitimate complaint that the bundles are unfair.

⁴² See, e.g., G.A. Cohen, *Rescuing Justice & Equality* (Cambridge, Mass: Harvard Univ. Press, 2008), pp. 93-94 and Serena Olsaretti, "Responsibility and the Consequences of Choice," *Proceedings of the Aristotelian Society* CIX (2009): 165-88.

Part III.

A. Given the examples that choice-sensitive egalitarians have used to motivate their view—examples in which choice and hard work coincide—I assume that most, if not all, choice-sensitive egalitarians agree with G.A. Cohen that the burdensomeness of an individual’s work matters to what they should be paid. In Cohen’s words, since “both money and quality of work experience matter to people ... both should therefore enter the egalitarian reckoning.”⁴³ In order to determine fair bundles, then, we need some way to judge how much additional income would fairly compensate for the various ways in which some work experience are worse than other work experiences. Cohen, however, punts on the question of how such judgments are to be made.⁴⁴ And, indeed, given the variety of ways in which a work experience can be worse, the task might seem to be unmanageable.⁴⁵ Yet, it turns out that we already have a device capable of outputting precisely these types of judgments. The following thought experiment both illustrates and makes precise the input that is relevant for determining fair bundles and how that input is to be used.

I begin with a brief overview and then describe the thought experiment in greater detail. Here is the brief overview: Imagine that the occupational positions that currently exist in the real world were auctioned off in a society in which all individuals were equally qualified for all occupational positions. The hypothetical auctioneer increases or decreases the income associated with each occupational position until the number of bidders for each occupational position equals the number of individuals in that occupational position in the real world. The auction continues until all participants declare themselves satisfied. The resulting income distribution reflects individuals’ preferences over occupational positions but not the unequal distribution of talents. Significantly, however, the hypothetical auction does not determine what individuals actually do. That depends on what talents individuals actually possess. Instead, the hypothetical auction determines how much occupational positions ought to be compensated (through taxation and redistribution) in the actual world.

Let me flesh out the thought experiment in more detail. The first step is to identify all of the occupational positions that currently exist in our society.⁴⁶ Occupational positions are not the same

⁴³ Cohen, *Rescuing Justice & Equality*, p. 200. See also *Ibid.*, p. 369.

⁴⁴ Cohen, “The Pareto Argument for Inequality,” p. 173, reprinted in *Rescuing Justice & Equality*.

⁴⁵ For example, some occupational positions are worse with respect to such factors as “difficulty, physical and mental challenge, monotony, labor schedule, degree of personal interaction, and transparent effectiveness.” Seana Valentine Shiffrin, “Incentives, Motives, and Talents,” *Philosophy & Public Affairs* (2010): 111-42, at p. 129. Other occupational positions are worse with respect to “dirtiness, stress, or seasonality of the job, risk of injury, lack of social status, etc.” Julian Lamont, “Incentive Income, Deserved Income and Economic Rents,” *The Journal of Political Philosophy* 5 (1997): 26-46, at p. 31. And these lists, of course, are not exhaustive.

⁴⁶ Strictly speaking, identifying the occupational positions that currently exist is not required. We could instead alter the types of occupational positions that exist in the real world. What is crucial, however, is that the occupational positions

as occupations. Rather, an occupational position refers to combinations of leisure and types of labor. For example, if an individual works full-time as a waitress and part-time as a cashier, then one of the occupational positions will be the combination of those two jobs. Once we have identified all the occupational positions, we must hold the number of individuals in each occupational position constant: if there are 300 corporate lawyers in the real world, then there will be 300 corporate lawyer positions available in the auction, and so forth.

The second step requires us to envision a society identical to our own except that everyone is equally qualified for all occupational positions. That is, everyone has the same natural talents and the same opportunities to develop them. Everything else, however, must remain constant. Specifically, individuals' preferences—and the strength of those preferences—over occupational positions must remain the same as in the real world. Thus, if you prefer the work involved in being a professor over the work involved in being a coal miner in the real world, you will continue to prefer the work involved in being a professor over the work involved in being a coal miner to the same extent in the hypothetical auction.

In the third step, we must imagine an auction in which the members of the society envisioned in step two bid for each of the occupational positions identified in step one. Starting from the default assumption of equal pay for all occupational positions, the hypothetical auctioneer increases or decreases the earnings associated with each occupational position until the number of bidders for each occupational position is equal to the number of individuals currently in that occupational position in the real world. The auction continues until each individual is assigned an occupational position and all participants declare themselves satisfied.

Since, by hypothesis, all of the participants in the auction are equally qualified for all occupational positions, the level of pay that would arise in the auction reflects the degree to which occupational positions are preferred or dispreferred. Occupational positions that are dispreferred will pay more; occupational positions that are preferred will pay less. To illustrate: since it is not particularly fun to be a coal miner, whereas many people enjoy playing sports, when the former position comes up in the auction, the auctioneer will need to set a high level of pay in order to entice a sufficient number of applicants. In contrast, when the job of basketball player comes up, even a relatively low level of pay is likely to attract plenty of applicants. Indeed, this confirms what

included in the auction be the same occupational positions that *will exist* in the real world. Otherwise, we will have succeeded in determining a fair distribution of bundles that do not exist. The easiest way to do this, of course, is to identify the occupational positions that currently exist. Nonetheless, the thought experiment leaves open the possibility of changing them. Indeed, the auction itself is a dynamic process in which occupational positions are likely to change: as the levels of pay change, certain occupational positions might cease to exist. For example, if nobody is super rich at auction levels of pay, occupational positions that cater to the super rich might disappear. Another auction must then be run with the new occupational positions.

we already knew to be true: the reason star basketball players earn so much more than coal miners in the real world is not because most people prefer coal mining to playing basketball, but rather because the natural talents necessary to be a star basketball player are extremely scarce.

To be slightly more accurate, the level of pay that would arise in the auction reflects the degree to which an occupational position and the necessary training for that occupational position are preferred or dispreferred. In the hypothetical auction, everyone is equally capable of developing the requisite skills for each and every job, but, in deciding how much income they would be willing to accept for any particular occupational position, they must factor in both the position itself and the necessary training. For example, in deciding what level of income they would accept to be a cardiologist, they must factor into their decision the years of medical training required.

In the fourth and final step, we use the levels of pay in the hypothetical auction to determine compensation for occupational positions in the real world. The goal of taxation and redistribution, then, should be to match the level of pay that would be reached in the hypothetical auction (absent other considerations, such as the need for incentives). Significantly, however, the thought experiment only determines the level of pay; it does not determine what individuals actually do.

B. The idea that an auction device can be used to determine a fair distribution is, of course, a familiar one. Dworkin, for example, uses an auction to divide resources fairly among castaways on a deserted island.⁴⁷ Similarly, Philippe Van Parijs flirts with the idea of using simultaneous auctions to determine employment rents (that is, the additional income individuals receive because jobs are scarce). Nonetheless, both Dworkin and Van Parijs ultimately reject the idea of using the auction device to determine fair pay. It is instructive, then, to see why.

Like my hypothetical equal endowment auction, Dworkin considers the possibility of assuming everybody is equally talented and then asking what jobs would pay in such a world. Yet, he immediately rejects the idea because, if the teenager next door plays basketball with equal skill as LeBron James, individuals would have little incentive to pay the ticket price to watch James play and, as a result, the occupational position of professional basketball player would probably cease to exist. In Dworkin's words: "in a world in which everyone could hit a high inside pitch or play sexy roles in films with equal authority, there would probably be no baseball or films ...".⁴⁸ The worry, then, is that we cannot use an auction to determine what jobs would pay in the absence of unequal endowments, because in such a world the type and distribution of jobs would change. That, of

⁴⁷ Dworkin, *Sovereign Virtue*, pp. 68-71.

⁴⁸ Dworkin, *Sovereign Virtue*, p. 92.

course, is true—unless we hold the type and distribution of occupational positions constant, as my auction does.⁴⁹

Van Parijs rejects the use of the auction device for a different reason. He begins with the observation that, in order to determine the rents, we must let both those who are qualified for the occupational position and those who are not qualified for the occupational position bid for the job. So far we agree: Everybody should be allowed to bid for all jobs. Van Parijs, however, goes on to worry that, if individuals are permitted to bid for jobs for which they lack the requisite talents, this could “lead to many jobs being held by people who like them a lot but are utterly incompetent at them”.⁵⁰ That, of course, is true if the auction device is used to determine not only the bundles themselves but also the assignment of bundles. Yet, the problem does not arise if the auction is used solely to determine the level of pay that fairly compensates for different occupational positions.⁵¹ The key, then, is to set the assignment function aside.

Since my dismissal of the assignment function might seem ad hoc, let me conclude my argument by showing that the assignment function is in fact dispensable. Assume, for the sake of argument, that we have determined fair bundles of work experiences and compensation. Helen chooses to work overtime (an occupational position considered less desirable) and thus the bundle she receives includes more pay than Ivan’s. The distribution is fair despite the fact that Helen receives more. Yet, what explains the fairness? Is it the fact that Helen *chose* that occupational position or is it the fact that her occupational position was considered less desirable and thus merited additional pay?

To answer that question, imagine that we find out that determinism, of whatever sort is the most damning to choice-sensitive egalitarians, is true. Egalitarians presumably would not now insist on strict equality of pay, but rather would continue to allow inequalities in pay that reflect fair bundles. After all, taking into account differences in the quality of work experience does not *create* an inequality by paying some jobs more, but rather *preserves* equality, all things considered. Thus, the very same bundles of work experience and compensation would be justified, with or without choice. Yet, if this is the case, then there is no additional role for choice to play in justifying

⁴⁹ Had he considered it, Dworkin would presumably have rejected the auction I propose because the levels of pay do not “reflect the cost or benefit to others of the choices people make.” Ibid., p. 89. I intend to take up this issue in a separate paper, but am unable to address it here.

⁵⁰ Van Parijs, *Real Freedom for All: What (If Anything) Can Justify Capitalism?* (Oxford: Oxford Univ. Press, 1995), p. 123.

⁵¹ Of course, one of the results of restricting the auction’s role to determining fair bundles is that the resulting distribution (in the actual world) will not be envy free: some people might prefer the job-income bundle someone else has. Yet, as I argue elsewhere, envy-freeness should never have been the test of fairness in the first place. “Fairness, Envy-Freeness, and Taxation” (in progress).

inequalities in pay. The justificatory work is being done by whatever justifies the bundles themselves.

Conclusion

The fair bundles justification, as instantiated by the hypothetical equal endowments auction, is capable of explaining why the inequality between Helen and Ivan is fair. The reason it is fair for Helen to receive more income than Ivan is because the work she does—involving longer hours—is dispreferred and additional income is needed to compensate. Thus, the fair bundles justification is capable of justifying the inequality between Helen and Ivan with which I started this essay.

At the same time, however, the fair bundles justification avoids some of the problems that have long plagued choice-sensitive egalitarians. The first of these objections to choice-sensitive egalitarianism challenges the existence of a sharp distinction between the chosen and the unchosen.⁵² An individual's willingness to work hard, for example, might be just as innate as the individual's natural talents.⁵³ The fair bundles justification avoids this objection because the individual's *choice* is not what matters. Instead, what matters is what the individual *does*. According to the second objection, depending on the truth of determinism, individuals might not make choices in the relevant sense. Thus, we need to determine the truth of determinism before we can determine whether or not an inequality is fair. To quote G.A. Cohen, choice-sensitive egalitarians are "up to [their] necks in the free will problem."⁵⁴ The fair bundles justification also avoids this objection. It does not matter whether free choice is possible because that is not what the fair bundles justification is tracking. Finally, according to the third objection, choice-sensitive egalitarianism has the following unattractive feature: individuals are compensated not in virtue of their *equality* to others, but in virtue of their *inferiority*.⁵⁵ Yet, basing compensation on inferiority—and especially requiring someone to prove their inferiority as a condition of compensation⁵⁶—seems at odds with the spirit of egalitarianism. The fair bundles justification, however, avoids this objection as well. In determining how much to redistribute to, say, a food service worker, it does not matter whether the food service worker could have been a corporate lawyer because we are not

⁵² See, e.g., Scheffler, "What Is Egalitarianism?" p. 18.

⁵³ See, e.g., John Rawls, *A Theory of Justice* (Cambridge, Mass: Harvard Univ. Press, 1971), p. 312.

⁵⁴ Cohen, "On the Currency of Egalitarian Justice," p. 934.

⁵⁵ Elizabeth Anderson, "What is the Point of Equality?" *Ethics* 109 (1999): 287-337, at p. 306.

⁵⁶ Under Dworkin's account, for example, in order to show that their lower income is not attributable to choices, individuals might be required to prove that they have below average physical or mental abilities as a condition of receiving compensation. Dworkin, *Sovereign Virtue*, pp. 101-02.

compensating her for any supposed lack of talent. Rather, we are simply compensating her for what she does.

To summarize, then, here are my main claims: In Part I of this essay, I showed that choice-sensitive egalitarians have long conflated two justifications for inequalities in pay. In Part II, I argued that, at most, choice can only justify the assignment of bundles; it cannot justify the bundles themselves. Finally, in Part III, I provided an alternative justification for inequalities in pay—an alternative justification that escapes the problems that have long plagued choice-sensitive egalitarianism. Significantly, however, my argument is not merely that there is a different justification, albeit one that avoids certain problems, for the same inequalities that egalitarians already endorse. Rather, if right, my argument has implications for which inequalities egalitarians should endorse. Since the two justifications sanction radically different inequalities, the life prospects of coal miners, home health aides, and food service workers (at least in the ideal world) are very much at stake.