

# The Paradox of Exploitation

## Abstract

The concept of exploitation brings many of our ordinary moral intuitions into conflict. Exploitation—or to use the commonly accepted ordinary language definition, taking unfair advantage—is often thought to be impermissible. In order to be permissible, transactions must not be unfair. The claim that engaging in mutually beneficial transactions is morally better than not transacting is also quite compelling. However, when combined with the claim that morally permissible transactions are better than impermissible transactions, these claims imply the counterintuitive claim that it is obligatory to engage in mutually beneficial transactions. In this paper I outline the conditions that comprise this ‘paradox of exploitation’ along with a solution that involves replacing one of the original conditions with a condition I call Weak Non-worseness. The solution captures the priority of our concerns about exploitation by making a concern for the fairness of a transaction subsidiary to a concern for the welfare of the would-be exploited.

Wordcount: 9957

Recently the concept of exploitation has received renewed attention. Much of the philosophical discussion has centred on debates about what makes a transaction fair.<sup>1</sup> This is an important question since exploitation is generally thought to involve unfairness in transaction. Most of the remaining debate surrounds questions about whether conditions in addition to fairness—such as ‘attitudinal deficiencies’ on the part of the exploiter—are also necessary for exploitation.<sup>2</sup> Both of these questions are concerned primarily with identifying the scope of exploitation, with *what exploitation is*. However, as G.A. Cohen writes, we need more than “a precise definition of what exploitation is... we want to know exactly why it is wrong.”<sup>3</sup> Not only has this latter task been largely neglected, but as I will show, providing an account of *why* exploitation is wrong is also much more difficult than it might first appear.

Attempts to identify what is wrong with exploitation are often frustrated because the concept brings many of our deeply held moral intuitions into conflict. Because they incorporate intuitions about the impermissibility of engaging in unfair transactions alongside a desire to improve the welfare of those involved in the transactions, accounts of exploitation’s wrongness include concerns typically attributed to both non-consequentialist and consequentialist moral frameworks. This combination makes it difficult to articulate which moral principles are

---

<sup>1</sup>See, for example, Steiner (1984), Wertheimer (1996), Vrousalis (2013), Goodin (1987), and Roemer (1982a).

<sup>2</sup>See, for example, Sample (2003), Goodin (1987), Vrousalis (2013), and Veneziani (2013).

<sup>3</sup>Cohen (1995: 144).

violated in cases of exploitation and what our response to certain cases ought to be.

In this article, I describe a paradox that arises when we attempt to account for normative intuitions about transactions that are Pareto improving and consensual, but nevertheless unfair. The paper proceeds as follows: in the first section I outline a basic tension between a concern for fairness and a concern for welfare in exploitative transactions. In the second I show that an attempt to incorporate both concerns leads to what I call the ‘paradox of exploitation’. In section three I evaluate the conditions that generate the paradox and outline possible resolutions. In section four I argue for one of these resolutions, which, I argue, best incorporates both a concern for fairness and a concern for welfare. The fifth section concludes.

## 1 The Initial Tension

When someone steals my wallet, they take it without consent and make me worse off than I was. Conversely, when I buy an apple from a market stall, the transaction is (ordinarily) consensual and mutually beneficial. Some accounts of exploitation include the first kind of interaction in the scope of exploitation. They allow exploitation to extend to non-consensual or harmful transactions. Other accounts limit their scope to consensual and mutually advantageous transactions.<sup>4</sup> I do not wish to take a stand on whether the latter forms of interaction are ‘really’ exploitations. However, the focus of this paper is on transactions commonly described as exploitative, but which are consensual and mutually beneficial. Such transactions are far more puzzling, partly because the following claim has some *prima facie* appeal:

**Pareto Permissibility.** If a transaction is consensual and strictly Pareto improving, then it is permissible.

In consensual and strictly Pareto improving transactions, both parties are better off having transacted and no other parties are burdened with negative externalities.<sup>5</sup> Further, when both parties *actually* consent to the transaction, concerns about paternalism and certain forms of force do not arise.<sup>6</sup> In what follows I will refer to Pareto improving, consensual, but nevertheless unfair transactions as exploitations.<sup>7</sup>

The Pareto Permissibility claim implies that some paradigmatic cases of exploitation are permissible. Consider a sweatshop labour contract. Here, perhaps more than anywhere else, the term ‘exploitation’ is used to condemn a practice that many argue is contemptible and morally amiss. Sweatshop labourers work long hours at near-subsistence wages in order to produce household goods that

---

<sup>4</sup>See Wertheimer (1996: 14) and Wertheimer (2007) for a discussion of this issue.

<sup>5</sup>Throughout this paper I will always assume the strict version of Pareto improvement.

<sup>6</sup>We can imagine a Pareto improving transaction to which Bob would have consented, but did not *actually* consent, perhaps because Carol made the transaction on Bob’s behalf (and without his permission) or because Alice forced Bob to transact. In both cases it seems that the lack of actual consent gives Bob a valid complaint against Carol (and Alice). However, note that for revealed preference theorists and subjective welfarists, strict Pareto implies hypothetical consent.

<sup>7</sup>This definition is purely conventional—as noted above, I do not want to take a stand here on issues of scope.

the workers cannot themselves afford to buy. Robert Ross’s account of conditions at Chentex, a Nicaraguan garment firm subcontracted by retail chain stores is, sadly, typical of this form of labour:

Chentex workers earn less than 1 percent of the retail price of the jeans they stitch—between thirty and forty cents an hour. This compares to the 10 percent typical of the global north.

When we visit the workers’ homes... we meet a woman, Cristina, who was fired from Chentex... Christina’s home is a wooden frame, ten-feet square, hung with plastic sheeting for two of the walls and with cardboard boxes that once held shirts shipped from the freezone in Panama for the rest. Her shack has a dirt floor and holds one large bed and (barely) two chairs for herself, her husband, and their baby. Her toilet is a hole in the ground.<sup>8</sup>

Christina’s condition is appalling and her wages—when she was employed—were low (around \$4 per day), but her story is tragically common.

Nevertheless, workers in Chentex factories often receive pay that exceeds wages offered elsewhere in Nicaragua. Ross notes that per capita income in Nicaragua “is about \$470 annually,” but garment workers can earn two or three times as much.<sup>9</sup> Consequently, jobs like Christina’s are in high demand. Not only do Chentex workers earn more than their compatriots, they also freely choose—in fact, compete—to work in sweatshop conditions. Christina’s employment at Chentex was both consensual and Pareto improving. Thus, according to Pareto Permissibility the transaction was morally permissible. Yet, most would resist this conclusion. Instead, they would draw the opposite conclusion, outlined in the following claim:

**Moral Impermissibility.** There exist some consensual and strictly Pareto improving transactions (e.g., sweatshop contracts) that are morally impermissible.

Pareto Permissibility and Moral Impermissibility obviously conflict. The latter is simply the negation of the former. It cannot be the case that all strictly Pareto improving and consensual transactions are permissible and that the transaction between Christina and Chentex is strictly Pareto improving, consensual, but nevertheless, impermissible.

## 1.1 Responding to the Tension

One response to the tension between Pareto Permissibility and Moral Impermissibility, which I will address briefly before setting aside, is to claim that there are no transactions that actually satisfy Moral Impermissibility; either because the consent of those in positions of severe disadvantage cannot be considered *true* consent, or because although they appear beneficial, exploitative transactions are not *really* Pareto improving. These claims are certainly true of some, and perhaps many cases. Many labour contracts that initially appear beneficial, may turn out to be harmful, or to involve negative externalities. However, it

---

<sup>8</sup>Ross (2004: 116–18).

<sup>9</sup>Ross (2004: 114).

is implausible that all sweatshop contracts ordinarily described as exploitative are, all things considered, harmful. A similar response can be given to those who deny that workers' consent is *true* consent when their choice conditions are constrained. In addition to being paternalistic (by denying Christina's own claims that she is better off working at Chentex), these responses are rather unimaginative. Even if no *actual* cases of sweatshop labour are consensual and strictly Pareto improving, it seems possible to imagine transactions that are.

I should note that this dismissal comes with two qualifications. The first is that I do not claim that determining whether Pareto and consent are satisfied in certain cases is unimportant. It may be that nearly all cases ordinarily called exploitations are either non-consensual or harmful and wrong for precisely these reasons. However, my primary interest in this paper is with the question of whether (and how) transactions that are mutually beneficial and consensual can be impermissible. The second is that I do not deny the consent given in cases of exploitation is, in some way, morally suspect. But we should distinguish a thicker notion of consent from a thinner, descriptive use. Many of those involved in transactions we call exploitations (such as Christina) clearly consent in a *thin* sense, but they do so grudgingly, in unjust conditions and thus we may say they do not consent in a thick sense. In what follows I use consent in a thin sense. I place consent in a thick sense under the category of unfairness. Or, more accurately, the background conditions of injustice under which a transaction is non-consensual in a thick sense (whatever these may be) also suffice to make the terms of the transaction unfair.

With these qualifications in place, consider the tension between Moral Impermissibility and Pareto Permissibility. A handful of commentators have (somewhat notoriously) dismissed Moral Impermissibility, concluding that mutually beneficial and consensual transactions simply are not impermissible. A startling example is Nicholas Kristof's claim that "the central challenge in the poorest countries is not that sweatshops exploit too many people, but that they don't exploit enough."<sup>10</sup> Similarly, Matt Zwolinski claims that "since ... providing no monetary benefit does not violate anyone's rights, and since a contract whereby sweatshops agree to provide some benefit does not in itself violate anyone's rights, it follows that such contracts are not exploitative."<sup>11</sup> Broadly speaking, arguments against Moral Impermissibility claim there are no other wrong-making properties that override Pareto and consent.

## 1.2 Fairness and Welfare

Yet, there seem to be *many* conditions beyond Pareto and consent that make transactions wrongful. Philosophers have variously argued that transactions are wrongful when they are disrespectful, when they involve trade in non-commodifiable goods, when they violate moral rights, and (most applicable to the present discussion) when they are unfair. Clearly we need not look far for reasons to reject Pareto Permissibility. When considering exploitation, a concern for fairness stands out as an obvious wrong-making property. Indeed, most

---

<sup>10</sup>Kristof (2009: 14D).

<sup>11</sup>Zwolinski (2007: 711). More generally, Pareto Permissibility may be attractive to those who endorse the welfarist thesis that the moral goodness of a state of affairs is an increasing function of the sum (or average) of the welfare in that state.

theories of exploitation begin with the claim that exploitation involves taking *unfair* advantage.<sup>12</sup> Consider the following in place of Pareto Permissibility:

**Necessary Fairness.** Transactions are permissible only if they are consensual, strictly Pareto improving, and *fair*.

Necessary Fairness adds a fairness condition alongside strict Pareto and consent and makes these three properties necessary, rather than sufficient, for the permissibility of a transaction (since there are, of course, many other reasons for claiming a particular transaction is impermissible). Necessary Fairness entails that transactions are impermissible when they are unfair. Though existing accounts of exploitation disagree about how the unfairness in exploitative transactions should be explicated, there is significant agreement that conditions beyond strict Pareto and consent—which, because they pertain to distributive concerns, fall broadly under the label ‘fairness’—must be satisfied in order to ensure the permissibility of transactions.<sup>13</sup> Unfairness is one reason for the intuition captured by Moral Impermissibility that some Pareto improving and consensual transactions, such as sweatshop labour, are nevertheless impermissible. Consequently, unfairness also provides a reason to reject Pareto Permissibility. Indeed, this appeal to fairness is so obvious that we may wonder what motivates the defenders of Pareto Permissibility.

One charitable response to Pareto Permissibility is that those drawn to the condition really intend to defend a closely related claim, similar to what Alan Wertheimer has called ‘non-worseness’:<sup>14</sup>

**Non-worseness.** If a transaction is strictly Pareto improving and consensual, then engaging in the transaction is better than not transacting.<sup>15</sup>

Although similar to Pareto Permissibility, Non-worseness does not conflict with Necessary Fairness. It allows that transactions can be strictly Pareto improving and consensual, yet impermissible because unfair. Non-worseness claims that, provided Pareto and consent are satisfied, transacting (though perhaps impermissible) is nevertheless better than not transacting. The endorsement of both Necessary Fairness and Non-worseness is an attractive resolution of the initial tension because this approach captures the idea that there is something good about Pareto improving, consensual transactions—they are better than not transacting—and yet it is consistent with the intuition that some of these transactions, though better than not transacting, may nevertheless be impermissible.

---

<sup>12</sup>See Elster (1982), Goodin (1987), Sample (2003), Steiner (1984), and Wertheimer (1996).

<sup>13</sup>For the purpose of the paradox that follows, it is unnecessary to specify a particular conception of fairness, though the resolution I defend in section four places some broad constraints fairness criteria.

<sup>14</sup>Wertheimer offers his own analysis of non-worseness in Wertheimer (1996: 289–293).

<sup>15</sup>Wertheimer states the claim as follows: “Given that I have a right not to transact with B and that transacting with B is not worse than not transacting with B, it can’t be seriously wrong for me to engage in an unfair transaction with B” (Wertheimer 1996: 289). Stated in this way, Wertheimer’s Non-worseness is not equivalent to the condition I outline above. The negation of ‘worse than’ is ‘better than or equal to thus, Non-worseness in Wertheimer’s use implies that Pareto and consent are sufficient for transacting to be better than or equal to not transacting. Throughout I will interpret ‘betterness’ as *moral* betterness, unless explicitly stated.

Indeed, both conditions lie at the heart of our thinking about cases of exploitation. Exploitation involves a fundamental concern for the welfare of the exploited, but this concern differs from a general concern that their welfare is low. It is a concern that the welfare of the exploited is not being increased *enough* in a transaction. Though a particular transaction results in a welfare gain for the exploited, the claim that they are exploited entails that they *ought to* receive more. Thus, our concerns for exploitation involve both a concern for fairness—that the gains from transacting are fairly distributed—and welfare, that people do better than they otherwise would have done. Unfortunately, though Necessary Fairness and Non-worseness are individually compelling, jointly satisfiable, and essential to our thinking about exploitation, when combined with two compelling additional claims they generate a paradox.

## 2 The Paradox

Suppose Bob wants to purchase an apple from Alice and that the most he is willing to pay for the apple is \$2.50. Also suppose Alice’s reservation price is \$1; at any price below \$1 Alice is better off keeping the apple than selling it.<sup>16</sup> Together these suppositions mean that any transaction that prices the apple above \$1 and below \$2.50 will be strictly Pareto improving. That is, at any price in that range both Alice and Bob will strictly prefer transacting to not transacting. Suppose also that the most plausible account of fair transaction (which, for the time I need not specify) entails that the fair price for the apple is \$1.50. In order to focus on the issue of fairness I will assume—both here and throughout—that the transaction does not violate any other possible restrictions on the permissibility of transactions, including (thin) consent. I will limit my attention to cases where fairness is the only ‘permissibility variable’ in play. Finally, suppose that Alice offers to sell Bob an apple for the *unfair* price of \$2.

Assuming Bob wants to minimise the price he pays for an apple, Bob will prefer the fair price of \$1.50 to the unfair price of \$2, since \$1.50 is less than \$2. Bob will also prefer the unfair price to not transacting, since the unfair price of \$2 remains below his reservation price of \$2.50.<sup>17</sup> Bob’s preference order is:

$$\textit{Fair} (\$1.50) \succ \textit{Unfair} (\$2) \succ \textit{No transaction} \quad (1)$$

Now consider Alice. Assuming Alice wants to maximise her profit, she will prefer the unfair price (\$2) to the fair price (\$1.50) since \$2 is greater than \$1.50. She also prefers the fair price to not transacting, since \$1.50 is above her reservation price. Thus, Alice’s preference order is:

$$\textit{Unfair} (\$2) \succ \textit{Fair} (\$1.50) \succ \textit{No transaction} \quad (2)$$

Although Alice and Bob have opposing preferences between fair and unfair prices, both prefer some form of transaction to no transaction. If both also consent to the transaction, then it seems that transacting at any price in the

<sup>16</sup>Alice’s reservation price may, but need not, be a function of the cost of producing the apple. If it is, then when Alice sells the apple for less than \$1 she makes a loss. But she may have other reasons for setting her reservation price at \$1.

<sup>17</sup>Thus, I am not assuming Bob’s (and Alice’s) self-regarding preferences include a preference to satisfy a moral obligation to transact fairly. I will return to this issue in section 3.3.

range between \$1 and \$2.50 is not only prudentially better than not transacting, but it is also morally better. That is, as Non-worseness claims, it is better to transact unfairly (at \$2) than to not transact, since strict Pareto and consent are satisfied. With a particular case before us it is difficult to argue against Non-worseness. After all, both Alice and Bob would prefer transacting at \$2 to not transacting, both consent to this transaction, and—by Pareto—there are no negative externalities involved in the transaction. In fact, it seems that the only harm would come from *preventing* the transaction. Those who would deny Non-worseness owe an account of why forbidding transaction in the name of some other value should override the transactors’ preferences and welfare. Absent an explanation, the denial of Non-worseness appears perniciously moralistic.

As I noted, Non-worseness captures concerns similar to those expressed by Pareto Permissibility, but unlike Pareto Permissibility it is consistent with Necessary Fairness. However, one difference between Non-worseness and Pareto Permissibility is that the ‘better than’ relation expressed in Non-worseness—the claim that transacting is *better than* not transacting when Pareto and consent are satisfied—is an *evaluative* relation that provides an ordinal ranking of options. The relation expressed in Necessary Fairness (and Pareto Permissibility) on the other hand is a *deontically modal* concept that partitions acts into the two jointly exhaustive and mutually exclusive categories of ‘permissible’ and ‘impermissible’.<sup>18</sup> While one act can be better than another, it cannot be more permissible than another permissible act. Unlike goodness, permissibility does not come in degrees.<sup>19</sup>

Relating evaluative and deontic judgements is difficult. An act’s permissibility does not, generally, give us any information about the act’s goodness, and vice versa. However, there is one plausible exception to this claim. It seems that permissible acts are better than impermissible acts.<sup>20</sup> That is,

**Betterness of Permissibility.** For any two acts  $x$  and  $y$ , if  $x$  is permissible and  $y$  is impermissible, then  $x$  is better than  $y$ .<sup>21</sup>

<sup>18</sup>Here, and throughout, I will assume permissible and impermissible are mutually exclusive and exhaustive.

<sup>19</sup>It might be argued that one can make sense of one act being ‘more permissible’ than another. Insofar as this claim makes sense, it seems that it is because it references the goodness of the ‘more permissible’ act. Martin Peterson (2013) has defended the notion of ‘degrees of rightness’, though his is a minority view.

<sup>20</sup>Some non-consequentialists may reject this claim. However, as I discuss in section 3.1 there are good reasons to ensure our normative theories respect it.

<sup>21</sup>Betterness of Permissibility (BP) is a stronger version of a related condition, Monotonicity of Impermissibility (MI), which states that acts worse than impermissible acts are also impermissible. BP implies MI, but MI does not imply BP.

Proof:

BP  $\rightarrow$  MI: if MI is violated, then so is BP. Suppose, violating MI, an act worse than an impermissible act is permissible: it is impermissible to  $x$ , permissible to  $y$  and  $y$  is morally worse than  $x$ . Then there is a permissible act,  $y$ , that is not better than an impermissible act,  $x$ .

$\neg$ (MI  $\rightarrow$  BP): if BP is violated, then MI can still be satisfied. BP is false in degenerative cases where all acts have equal value, such as when  $x$  is permissible,  $y$  is impermissible and  $x$  is equally as good as  $y$ . But here MI is not violated, since MI says nothing about the impermissibility of acts that are equally as good as an impermissible act.  $\square$

The Paradox can be generated by appealing to the weaker MI; however, the ease of interpretation of BP make the stronger condition much more straightforward. Also note that, though the ordering in (3) is inessential for the paradox, it depends on the stronger BP condition.

For example, suppose that theft is impermissible, and murder is worse than theft. Betterness of Permissibility implies that murder is also impermissible. Impermissible acts do not suddenly ‘flip’ and become permissible as the goodness of the acts in question diminish.

Recall the earlier example of the apple sale. By hypothesis, the price of \$1.50 for the apple is fair and consensual, though all prices above \$1 and below \$2.50 are strictly Pareto improving. Since the transfer of the apple for \$1.50 is fair, consensual, strictly Pareto improving (and satisfies other possible conditions), it is permissible. However, because the \$2 price does not satisfy the fairness condition it is impermissible, according to the Necessary Fairness premise. If permissible acts are better than impermissible acts, then selling the apple for \$1.50 is better than selling the apple for \$2. That is,

$$\textit{Fair} (\$1.50) > \textit{Unfair} (\$2) \quad (3)$$

In addition, Non-worseness claims that if a transaction is strictly Pareto improving and consensual, then engaging in the transaction is better than not transacting. Since transacting at \$2 satisfies these conditions, transacting at \$2 is better than not transacting. So, by Non-worseness,

$$\textit{Unfair} (\$2) > \textit{No transaction} \quad (4)$$

The combination of these two orderings gives us the following goodness ordering for the three options:

$$\textit{Fair} (\$1.50) > \textit{Unfair} (\$2) > \textit{No transaction} \quad (5)$$

Ordering (5) can explain the protracted disagreement between those who would condemn sweatshops on grounds of exploitation and those, like Kristof, who argue that we need more sweatshops. In a sense, both are correct. Kristof, who argues it is better for Christina to be employed by Chentex rather than unemployed can endorse Non-worseness, which gives ordering 4. Those who condemn sweatshops can endorse the Necessary Fairness condition, which—along with the Betterness of Permissibility—gives ordering (3). Both parties can endorse ordering (5). While employment is better than unemployment, employment at a fair wage is better still.

We now have two attractive results. First, as I noted above, Non-worseness is compatible with Necessary Fairness, allowing for the claim that there is *something* good about Pareto improving, consensual transactions, alongside the claim that these transactions may nevertheless be impermissible. Second, Non-worseness, Necessary Fairness, and Betterness of Permissibility resolve the disagreement between sweatshop defenders and critics: though it is true that sweatshops are better than nothing, they are not better than situations where the terms of the labour contract are fair. While Kristof is correct that we need more labour contracts to lift those in developing economies out of poverty, these contracts should also involve a fair distribution of the gains from transacting. Thus, the three conditions not only capture the concerns about fairness and welfare present in our thinking about exploitation, they also resolve the debate that motivates the initial tension I described in the first section.

Unfortunately, although these two results are attractive, a problem looms. Taken together, Non-worseness, Necessary Fairness, and Betterness of Permissibility conflict with the following claim:



**The Optionality of Ordinary Transactions.** Strictly Pareto improving and consensual transactions are not *ordinarily* obligatory.

The ‘ordinarily’ clause in the condition restricts its scope to day-to-day transactions that do not feature extra-ordinary moral obligations. For example, Alice may have an obligation to purchase Bob’s car if she promised to do so. Promises, reparations, duress, and similar states can make certain transactions obligatory. The Optionality premise does not deny these facts. Rather, it claims that in ordinary cases, like buying an apple from Alice, we are not obliged to transact even if the transaction would be mutually beneficial and consensual. In other words, the optionality condition denies that transactions become obligatory simply in virtue of being strictly Pareto improving.<sup>22</sup>

The conflict between the Optionality premise and the conjunction of Non-worseness, Necessary Fairness, and the Betterness of Permissibility occurs in the following way. According to Necessary Fairness, fairness is necessary for the permissibility of a transaction. In the fruit sale case, because the \$2 transaction is (by hypothesis) unfair it is impermissible. Non-worseness claims that if a transaction is strictly Pareto improving and consensual, it is better than not transacting. Thus, as expressed in ordering (4), not transacting is worse than engaging in the \$2 transaction. Betterness of Permissibility implies that if an act is worse than an impermissible act, then it too is impermissible. Since (by Non-worseness) not transacting is worse than transacting unfairly, and (by Necessary Fairness) transacting unfairly is impermissible, then (by Betterness of Permissibility) not transacting is also impermissible. Yet, this result conflicts with the Optionality premise.<sup>23</sup> Surely Bob does nothing wrong when he fails to purchase an apple from Alice. Either the Optionality premise, or one of the three premises entailing the *obligation* of ordinary transactions must be abandoned.<sup>24</sup>

<sup>22</sup>It is important to note that if Bob fails to engage in a Pareto improving and consensual transaction with Alice he is *irrational*. While the claim that the set of moral acts must be a subset of rational acts is compelling, the claim that the two sets are identical seems implausibly strong. Indeed, it is likely that many forms of consumer behaviour are irrational, but not impermissible. See Tversky (1981).

<sup>23</sup>Here I am assuming that  $(\Diamond(x \wedge \neg x) = \neg \Box x)$ . That is, the Optionality condition’s claim that it is optional to transact implies that it is not obligatory to transact, which contradicts the implication of the first three conditions that it is obligatory to transact.

<sup>24</sup>It can be shown that the paradox also arises when Wertheimer’s form of Non-worseness (Pareto and consent are sufficient for transacting to be better than *or equal to* not transacting) replaces the version defined above and when Monotonicity of Impermissibility (MI) replaces Betterness of Permissibility. However if *both* alternative conditions are used, then these conditions along with Necessary Fairness do not contradict the Optionality condition.

If transacting unfairly at \$2 is impermissible, then this is, by Wertheimer’s Non-worseness condition, morally better than or equal to not transacting. If it is equivalent, then MI does not imply that not transacting is impermissible. MI only states that an act that is *worse* than an impermissible act is also impermissible. Thus, when both weaker conditions are adopted there is no paradox. But this is an unattractive solution. The paradox can be reinstated with the addition of the rather weak claim that there can be no difference in permissibility without a difference in moral value. (I discuss the problematic consequences of denying this claim in the following section.) With this added condition, if transacting unfairly at \$2 is impermissible and not transacting is morally equivalent in value to transacting at \$2, then not transacting is also impermissible, contradicting the Optionality condition.

### 3 Evaluating the Conditions

Non-worseness captures the common consequentialist concern to promote individual welfare. Necessary Fairness captures the non-consequentialist idea that those who transact unfairly do something impermissible. These conditions—which are compatible—capture concerns for welfare and fairness that shape our thinking about exploitative transactions. Not only do they capture the moral concerns that feature in exploitation, as we saw in the previous section, they also help explain the disagreement between defenders and critics of sweatshops.

The remaining two conditions—Betterness of Permissibility and Optionality of Ordinary Transactions—can be interpreted as conditions that respectively prohibit undesirable features of non-consequentialist and consequentialist moral frameworks. Optionality says that engaging in Pareto improving and consensual transactions should not be morally required. If we endorse both Optionality and Non-worseness, then this is true even if doing so is morally better. In this sense, Optionality can be seen as a condition that allows for supererogation. Or put in another way, Optionality rules out a form of the demandingness challenge to consequentialism, which claims the obligation to perform the act that brings about the most good is too demanding.

#### 3.1 Betterness of Permissibility

The theoretical appeal of Betterness of Permissibility is less obvious. The condition specifies a relation between evaluative claims about the goodness of a state of affairs and deontic claims about the rightness of acts. Such ‘bridge principles’ are controversial and, indeed, the Betterness condition is violated within many non-consequentialist moral frameworks. However, Betterness of Permissibility can be interpreted as a principle that prohibits the paradox of deontology, and thus, may be interpreted as a consistency requirement for a moral framework.

The connection here can be illustrated by considering the two cases in which Betterness of Permissibility is violated. First, note that in order for Betterness of Permissibility to be directly violated within a moral framework—rather than merely inexpressible—the framework must make claims about both rightness and goodness.<sup>25</sup> The ability to explain not only what acts are permissible or impermissible, but also an explanation of why some acts are worse (or better) than others is a desideratum of any complete moral framework. A moral theory that cannot express moral differences between theft and murder, or murder and genocide is woefully incomplete.

A moral framework violates Betterness of Permissibility whenever an act lacks the property (or properties) that, according to the theory, makes an act permissible, but has more of the property (or properties) that, according to the theory, make acts morally good. In other words, Betterness of Permissibility is false when there is an impermissible act that is better than a permissible act. Though violations of the condition are possible for frameworks that include both accounts of goodness and permissibility, such violations encounter a dilemma: If a moral framework violates Betterness of Permissibility, then it is either ‘axiologically incongruous’ or it encounters the ‘paradox of deontology’.

---

<sup>25</sup>For example, some scalar consequentialists, such as Alastair Norcross (2006) argue we should abandon the notion of rightness.

Consider the first horn. If for a particular framework the properties that determine the goodness of a state are wholly independent from the properties that determine the rightness of an act, then within that framework doing the right thing will have nothing to do with what—*according to that same framework*—makes the world a better place. If the properties that determine rightness and goodness are unrelated then the moral framework suffers from a kind of ‘two-mindedness’ about moral value, it is, in this sense, axiologically incongruous. Further, if goodness and rightness are not connected, then right acts cannot be justified by appealing to their *value*. Yet, without such a connection it seems very hard to explain or justify an account of rightness to others who do not accept it. This fork of the dilemma should be resisted.

In order to avoid this problem, the properties that determine rightness should be tied in some way to the properties that determine goodness. That is, a reason impermissible acts are prohibited is that they make the world a worse place. This leads us to the second horn. If the properties that determine rightness and goodness are tied together in this way, then it seems inconsistent to claim that a concern for rightness justifies prohibiting act  $\psi$  if the performance of act  $\psi$  will bring about more of the property that the concern for rightness is intended to protect or promote. In other words, why not violate a restriction whose point it is to promote a property if doing so leads to *more* of that property in the world? Principles that prohibit what they are designed to promote suffer from a form of inconsistency. This problem is one form of the well known paradox of deontology.<sup>26</sup>

Note that there is a second way Betterness of Permissibility can be violated. The condition is false when a permissible act and an impermissible act are equally good. Here again we face a dilemma. Suppose goodness and rightness are related. Then in this case, the two cannot be equivalent, for if they differ in rightness and rightness contributes to goodness, then they cannot be equally good. The situation is definitionally impossible. Suppose that rightness does not contribute to goodness. Then, though there is no definitional inconsistency, we again face axiological incongruity: why we should care about the right making property since there is, by supposition, no moral value for the property.

In response to the dilemmas that arise from the violations of Betterness of Permissibility, non-consequentialists may claim that what makes an act good is part—but not all—of what makes the act right. That is, the right-making properties are part of, but not identical to, the good-making properties. In general, there may be many different properties that are relevant for rightness and many different properties relevant for goodness. However, these mixed accounts do not avoid the dilemma. For any property apart from goodness that could be relevant to rightness we face the question, ‘why does this property matter for rightness?’ Any answer that says this property is right-making, but not associated with goodness is axiologically incongruous. And any answer that ties the property to goodness faces the paradox of deontology if it violates the Betterness condition in the first way and cannot arise if it violates the condition in the second way. Thus, the dilemma plagues both ‘pure’ and ‘mixed’ accounts.

If a moral framework violates the Betterness of Permissibility by claiming impermissible acts are better than permissible acts, then it faces either the paradox of deontology or it is axiologically incongruous. If a moral framework violates

---

<sup>26</sup>See, for example, Nozick (1974), Scheffler (1988), and Heuer (2011).

the condition by claiming a permissible act and an impermissible act are equally good, then the two conditions must be unrelated and, again the framework is axiologically incongruous. Finally, any ‘mixed’ account of the relationship between goodness and rightness can be decomposed into a ‘pure’ account where these problems re-emerge. Assuming non-consequentialists would rather confront the paradox of deontology than axiological incongruity, we may interpret the Betterness of Permissibility as a consistency requirement that avoids the paradox of deontology.

In summary, the first two conditions, Non-worseness and Necessary Fairness respectively capture consequentialist and non-consequentialist concerns present in our thinking about exploitation and they can explain ongoing disagreement about cases of exploitation. The Optionality and Betterness conditions can be seen as intuitions that our analysis of exploitation should be able to account for supererogation and avoid the paradox of deontology. Thus, the four conditions that comprise the paradox ask us to capture both consequentialist and non-consequentialist concerns while avoiding the well-known problems that confront each moral framework. Seen in this light, the fact that these four conditions are jointly inconsistent is perhaps unsurprising.

### 3.2 Two Traditional Responses

For traditional consequentialists and non-consequentialists alike, the inconsistency is not only unsurprising, but a number of resolutions readily suggest themselves. Consequentialists have developed responses that diminish the force of demandingness objections.<sup>27</sup> Armed with these responses they may find ways to bite the Optionality bullet, while retaining the remaining conditions. Alternatively, consequentialists may reject the Necessary Fairness condition, claiming that fairness is valuable only insofar as it contributes to moral goodness. They may acknowledge that fairness is a good-making feature, but deny that it is strictly necessary for permissible transaction. Similarly, Non-consequentialists may be untroubled by the paradox of deontology. They may grant that it represents a challenge to their framework, but argue that the theoretical virtues of non-consequentialism—for example, its close connection to ordinary morality and its accommodation of supererogation—outweigh this cost. In this way Non-consequentialists might reject the Betterness of Permissibility. Alternatively, they may grant that Pareto improving and consensual transactions are indeed *prudentially* better than not transacting, but deny these features entail that transactions are truly *morally* better.

These are reasonable responses. The four conditions that comprise the paradox are all intuitively compelling. Consequently, solutions that reject one condition and retain three are also likely to be compelling, simply because they capture three reasonable claims. Broader theoretical commitments may provide reasons for rejecting one or another of the four conditions. Both ‘pure’ consequentialist and ‘pure’ non-consequentialist approaches to the tension I have identified above simply dissolve the problem that exploitation presents by fiat. From both perspectives, explaining what we should do in cases ordinarily described as exploitative is relatively straightforward. Though these responses are reasonable, a better response to the paradox would capture the spirit of all four

---

<sup>27</sup>See Sobel (2007) for such an account.

conditions. By fully abandoning any one of the conditions we fail to capture a plausible part of our ordinary thinking about exploitation, which incorporates concerns from both frameworks. In section four I outline a solution to the paradox that attempts to balance a concern for welfare against a concern for fairness and in so doing, retains the spirit of all four conditions. However, before doing so, I want to briefly consider the consequentialist challenge to Necessary Fairness since I believe this response is the most compelling alternative to the solution outlined in section four.

### 3.3 Revisiting Necessary Fairness

The consequentialist solution involves omitting Necessary Fairness and writing a concern for fairness into the transactors' preferences. The effect of this omission on the Paradox is relatively straightforward. If by Non-worseness unfair, Pareto improving transactions are better than not transacting and by Optionality not transacting is permissible, then by Betterness of Permissibility, transacting unfairly is also permissible. And of course, this conflicts with Necessary Fairness. If Necessary Fairness is abandoned, the conflict disappears.

When the concern for fairness is 'written into' individual preferences, rather than captured by Necessary Fairness, the goodness of a state of affairs can be determined solely by reference to the agents' preference orderings. In this way, the transactors can assign their own balance to fairness and (material) welfare. If they care a great deal about fairness, then they may decide that, though the transaction involves *material* gain, they want to refrain from transacting because the distribution is unfair; if they do not attach much weight to fairness, then they will simply prefer those options that maximise their material gain. Allowing the transactors themselves to weigh fairness and material gain eliminates the need for an external—and arguably paternalistic—ranking of the concerns.

Those who prefer this response to the paradox are likely motivated by the welfarism, the claim that the moral goodness of a state of affairs is an increasing function of the sum, or average, of the welfare in that state.<sup>28</sup> There exist many accounts of what welfare consists in—and thus many versions of welfarism—but perhaps the most common approach, which I will assume here, is to identify an individual's welfare with the extent to which his or her (rational, informed) preferences are satisfied.<sup>29</sup> Since welfarists associate the relative goodness of a state with the degree of satisfaction of individual preferences, the relative goodness of fair and unfair transactions will depend on how much individuals *themselves* prefer these states. Since these preferences will vary between individuals, welfarism does not, *independently* of these preferences, imply anything about the relative value of fair or unfair transactions. The goodness of fair and unfair transactions is non-comparable. Welfarism does imply that mutually beneficial transactions are morally better than non-transaction. Thus, welfarism generates the following goodness ordering:<sup>30</sup>

$$\text{Fair } (\$1.50) \parallel \text{Unfair } (\$2) > \text{No transaction} \quad (6)$$

<sup>28</sup>Sen (1979)

<sup>29</sup>There are objections to this interpretation of welfare—see Hausman (1995)—but neither the consequentialist solution nor my rebuttal depend upon this account of welfare.

<sup>30</sup>I use  $x \parallel y$  to express the non-comperability of the goodness ( $\neg(x \geq y)$  and  $\neg(y \geq x)$ ) of  $x$  and  $y$ .

We can say a little more about this goodness ordering for our transactors Alice and Bob by considering the four possible configurations of preferences. Each of the two transactors may prefer fairness to material gain, or material gain to fairness, thus there are four possible alignments of preferences.

Suppose a transactor attaches lexical priority to fairness and considers material gain a ‘tie-breaker’. In this case, both transacting fairly and not transacting will be preferred to transacting unfairly. A secondary concern for material gain means that fair transaction will be preferred to non-transaction. If we return to the apple transaction between Alice and Bob, we find that this overriding preference for fairness implies a preference ordering that coincides with goodness ordering (7). Now suppose a transactor attaches lexical priority to material gain, with fairness operating as a secondary concern. This generates preference order (2) for Alice and (1) for Bob, as we saw in section one. These alternatives are presented together in the following table (the corresponding ordering numbers are listed in brackets):

	Alice	Bob
Overriding Fairness	$\$1.50 \succ N \succ \$2$ (7)	$\$1.50 \succ N \succ \$2$ (7)
Overriding Welfare	$\$2 \succ \$1.50 \succ N$ (2)	$\$1.50 \succ \$2 \succ N$ (1)

When both have an overriding preference for fairness (the top row) their preference orders coincide and the outcome will be a (fair) transaction at \$1.50. When Alice, but not Bob, has an overriding preference for fairness they will also transact at \$1.50. Though they prefer \$1.50 for different reasons, their most preferred options nevertheless coincide. And when Bob, but not Alice, has an overriding preference for fairness they will transact at \$1.50 since this is the only point at which their reservation prices overlap. In each case, if the fair transaction does not occur (for whatever reason), then no transaction will occur because in each possible configuration of preferences, one of the transactors’ second most preferred options is to not transact. Thus, if either Alice and Bob have an overriding concern for fairness they will transact fairly, or failing this, not transact. Unfair transaction cannot occur in these three configurations. Finally, if *neither* have an overriding concern for fairness, then both will prefer some material gain to none and the outcome may be a fair or unfair transaction. The result will simply depend upon how they bargain.

Thus, the only case in which an unfair transaction can occur is one in which neither transactor has an overriding concern for fairness. Though both may attach some weight to fairness in this last case, they will nevertheless care *more* about maximising their material gain. If neither transactor has a strong preference for fairness, what reason do we have oppose this solution?

There is, I believe, a very good reason to reject both this consequentialist solution and, more generally, the welfarist theory that motivates it. Above we saw four possible ways two rational transactors’ could weigh their concerns for fairness and material gain. Empirical results from in behavioural economics are initially encouraging. They suggest that in some contexts, some relatively wealthy people do have an overriding concern for fairness.<sup>31</sup> However, there is

<sup>31</sup>This behaviour is observed, for example, in players’ behaviour in the ‘ultimatum’ game.

also an important group of individuals who clearly lack an overriding concern for fairness: the very poor. When persons' welfare is very low they are likely to attach a much greater weight to whatever material gains they can acquire. Material gain is simply a name for an increase in a large kind of goods, which, much more than abstract or moral goods, is subject to decreasing marginal value. That is, as our level of material possessions increases, the value of each additional material possession decreases. Consequently, if we have no or few material possessions, material gain is likely to outweigh a gain in non-material goods. And this means that the very poor will rarely, if ever, express an overriding concern for fairness over material gain. But if we follow welfarists in claiming that moral goodness increases with welfare (or preference satisfaction), then we cannot condemn unfair transactions with the very poor. This is appalling—when someone is in such a bad state that they lack the ability to protest (by refusing to transact) we should not conclude from their overriding preference for material gain that any Pareto improving and consensual transaction is permissible. In many cases the preference for material gain expressed by the poor will have been produced through injustice: had they a just amount of wealth they would likely express a preference for fairness.<sup>32</sup> The morally repugnant consequence of the welfarist solution is that those looking to maximally further their own interest while avoiding wrongdoing should engage in unfair transactions with the poor. Since the welfare that results from such transactions is greater than the welfare they enjoy without transaction, the transaction is morally permissible.

## 4 Resolving The Paradox

If you accept my arguments in favour of retaining Betterness of Permissibility and Necessary Fairness and want to avoid biting the bullet by denying the Optionality condition, then only one option remains: Non-worseness. In the sections that follow I will argue that the paradox of exploitation should be resolved through a weakening of the Non-worseness condition, combined with a clarification of what, precisely, it means for a transaction to be unfair.

### 4.1 The Resolution

Recall the original Non-worseness condition: 'if a transaction is strictly Pareto improving and consensual, then transacting is better than not transacting'. I argue that Non-worseness is compelling only when weakened:

**Weak Non-worseness.** If one or more strictly Pareto improving and consensual transactions are possible, then for at least one of these transactions, transacting is morally better than not transacting.

Consider the following goodness order:

$$\textit{Fair} (\$1.50) > \textit{No transaction} > \textit{Unfair} (\$2) \quad (7)$$

---

Oosterbeek (et. al.).

<sup>32</sup>Of course, welfarists will not deny that that other transactions may be better still. But provided they accept the Optionality and Betterness claims and reject Necessary Fairness, they are committed to Pareto permissibility.

This order violates the original Non-worseness condition because the unfair transaction, though Pareto improving and consensual, is ranked below non-transaction. However, (7) respects Weak Non-worseness since, of the two possible Pareto improving and consensual transactions, at least one is morally better than not transacting.

Not only does (7) satisfy Weak Non-worseness, it is also consistent with the remaining three conditions. By Necessary Fairness, unfair transactions are impermissible and by the Optionality of Ordinary transactions, not transacting is permissible. In (7) not transacting is morally better than unfair transaction. Because both permissible acts are ranked above the impermissible act of unfair transaction, order (7) also respects the Betterness of Permissibility. Therefore, when the original Non-worseness condition is replaced by Weak Non-worseness, the new condition is jointly consistent with Necessary Fairness, Betterness of Permissibility, and the Optionality of Ordinary Transactions.

## 4.2 A Complication

Though Weak Non-worseness is consistent with the remaining three conditions in order (7) there is a special class of cases in which it *appears* a conflict might still arise: when *any* possible Pareto improving transaction is unfair. Suppose there are two (and only two) possible Pareto improving transactions in the apple sale case, both of which are unfair. For example, suppose the fair price for the apple is again \$1.50, but the only possible strictly Pareto improving transactions between Alice and Bob are an apple for (the unfair price of) \$2 or an apple for (the unfair price of) \$1.75. According to Weak Non-worseness *at least one* of these transactions must be better than not transacting. Perhaps it is \$1.75, the price that is closest to the fair \$1.50 transaction. In this case, by Weak Non-worseness, transacting at \$1.75 is better than not transacting:

$$Unfair (\$1.75) > No\ transaction \quad (8)$$

But (8) violates Betterness of Permissibility. By Necessary Fairness transacting at \$1.75 is impermissible because it is unfair, which means that an impermissible act is better than a permissible act. Does this mean that Weak Non-worseness fails to solve the paradox in all cases?

No, it does not. At least some point in the interval between Alice's reservation price and Bob's reservation price must be fair. That is, if any strictly Pareto improving transaction is possible, then at least one such transaction must also be fair. If one of \$1.75 or \$2 is fair, then, in the above example Weak Non-worseness does not lead to a conflict with Betterness of Permissibility. In general, this means that the special class of cases in which Weak Non-worseness fails cannot arise.

But why should we believe that if there is any possible, and permissible, strictly Pareto improving transaction, then at least one such transaction must be fair? In order to see that it is, we must distinguish fairness as a distributive concept from fairness as procedural concept. The kind of unfairness we are concerned with when we say that exploitation involves taking unfair advantage is distributive unfairness; that is, in exploitations the distribution of gains from transacting is a *maldistribution*. These gains ought to have been distributed in some other way. What a person gains from transacting is the difference



between what they get from engaging in the transaction from what they get if they refrained from transacting. If Bob would have paid up to \$2.50 for the apple, then when he pays anything less than \$2.50, he gains by transacting. Similarly, if Alice would have sold the apple for any price greater than \$1, then when she receives anything more than \$1, she gains from transacting. When Alice and Bob transact with each other, they both gain at any price above \$1 and below \$2.50. What we are concerned with in cases of exploitation is how this mutual gain should be fairly divided—with how the cake should be cut. In this case, the fair division must be some division of the cake itself, that is, of the gains from transacting. Thus, a fair distribution in the context of transacting must lie somewhere in the interval between each individual's reservation price.

Transactions may, of course, be unfair for other reasons. For example, suppose a transaction in which Alice and Bob exchange an apple for \$1.50 *distributively fair*, but suppose also that Bob cut in front of Carol in the queue and bought the last apple, which Carol intended to buy. Here we may also describe the transaction as 'unfair', not because the distribution of the gains between the buyer and seller was a maldistribution, but because one buyer was unfairly excluded from the market. The transaction is *procedurally* unfair.

These two forms of unfairness are often linked—in many cases procedural unfairness may cause distributive unfairness. Suppose Carol would have paid \$2 for Alice's apple, but because she was prohibited from entering the market, Alice only received \$1 for her apple from Bob. In this case, both forms of unfairness are present. Carol's being barred from the market is procedurally unfair and that leads to an unfair distribution of the gains from transaction between Alice and Bob. However, when we claim that Nicaraguans are exploited by Chentex this is not a claim about *procedural* unfairness. It is a claim about distributive unfairness, a claim that they ought to receive more for their labour.

Yet, even if we restrict our attention to distributive fairness, it seems that there can be cases where the fair transaction is not found in the interval between the transactors' reservation prices. Suppose Alice steals Bob's car and offers to sell it back to him. If Bob has no choice but to pay Alice to get the car back (and presuming he wants to replace it), then he will be willing to pay up to the amount that replacing the car would cost—if a new car costs \$500 (and if Bob attaches no sentimental value to the car), Bob would be willing to pay Alice up to \$499. Suppose that Alice doesn't need to use the car herself, and so doesn't attach much use-value to the vehicle, but she does know she can sell the car elsewhere for \$100 and so will sell the car to Bob for anything above this amount. If it is true that the distributively fair price must lie between the transactors' reservation prices, then the fair price for the car must be between \$101 and \$499. Surely this is incorrect. The fair 'price' for the car is \$0.

The act that leads to the morally best state of affairs is one in which Alice simply returns Bob's car. Since Bob values the car more than any price below \$500, then it seems that the next best outcomes are those in which Bob gives amounts less than \$500 to Alice in exchange for his car (e.g. \$1 for the car is better than \$2, is better than \$3, etc.). And finally the morally worst ranked outcome is one in which there is no transaction (here Bob's loss is equivalent to a loss of \$500).

$$\text{Return car} > \text{Transactions } \$1 - \$499 > \text{No transaction} \quad (9)$$

Ordering (9) captures our intuitions about the goodness of alternatives in the

car case. Here the fair transaction is the car's free return. Because transactions \$1–\$499 are unfair, they are impermissible by Necessary Fairness. Weak Non-worseness is satisfied since at least one Pareto improving and consensual transaction (transactions \$1–\$499) is morally better than not transacting. But as before, Betterness of Permissibility implies that not transacting is impermissible, contradicting the Optionality of Ordinary transactions. However, in this case the optionality condition is not compelling. Transacting is not morally optional because Alice is *obliged* to engage in the only fair transaction—she is obliged to return the car.

There is an important difference between the car case and the apple case. In the car case, a prior injustice—the car's theft—makes transacting obligatory and, therefore, the reason the Optionality of Ordinary transactions is implausible in this context is that the 'ordinary' clause is not satisfied. The transaction is not an 'ordinary' transaction because Alice is obliged to pay reparations for the theft. The car case lies outside the domain of transactions in which the paradox arises. This means that my prior claim that the fair transaction must lie in the interval between the transactors' reservation prices is true only if there are no extra-ordinary obligations to transact.

'But,' it might be objected, 'aren't some exploitations like this—in some cases aren't those who are poor being sold something they are entitled to receive for free, as reparation for past injustice?' The answer is yes, some transactions ordinarily described as exploitations fall into this category. And in these cases, the fair and obligatory transaction will be one in which the exploited simply receive what they are owed without exchanging anything in return. The fair transaction will be *unilateral* and thus, not a Pareto improvement. But the question that the paradox poses is whether *absent* any extra-ordinary obligation, transactions become obligatory in virtue of being Pareto improving. The original Non-worseness, Necessary Fairness, and Betterness of Permissibility conditions imply that they do, contradicting the ordinary intuition, captured by Optionality of Ordinary Transactions, that they do not.

'Surely,' the objection continues, 'if we limit our conception of fairness to distributive fairness *and* we limit our domain to those in which there is no obligation to engage in a non-Pareto improving transaction, then the claim that the fair transaction must be some *distribution* of the gains from transacting—that it must be found in the interval between the transactors' reservation prices—is true by definition.' Yes. But this is precisely the point. Because exploitation is a matter of distributive unfairness (which may, nevertheless, be caused by procedural unfairness) and because we are attempting to determine whether Pareto itself can entail an obligation to transact, we must limit the conception and domain of fairness in this way. And given these restrictions, it follows, as I claimed above, that if there is any possible, and permissible, strictly Pareto improving transaction, then at least one such transaction must be fair.

So, to return to the issue that motivated this discussion of fairness, if the fair transaction must be a strictly Pareto improving transaction—if it must fall between the transactors' reservation prices—then this means that it is not possible for cases like that depicted in ordering (9) to arise. If such orderings cannot arise, then there can be no case in which Weak Non-worseness fails to solve the paradox. Therefore, Weak Non-worseness, Necessary Fairness, Betterness of Permissibility *and* the Optionality of Ordinary Transactions can be jointly satisfied.

### 4.3 A Justification

Although Weak Non-worseness succeeds in leaving many of our intuitions about exploitation intact, without an independent justification for the condition, it feels *ad hoc*. Luckily, such a justification is available.

In the first section I claimed the fundamental concern in cases of exploitation is the welfare of the exploited is not being increased *enough* in a transaction. This concern for welfare is not allayed by morally obliging situations where no strictly Pareto improving transaction takes place, for here the would-be exploited have the lowest welfare. As Joan Robinson famously claimed, “the misery of being exploited . . . is nothing compared to the misery of not being exploited at all.”<sup>33</sup> The priority of a concern for welfare means that any criteria of fairness must morally permit a possibility for the would be exploited to permissibly engage in a transaction in which they better themselves. In other words, as I explained above, absent extra-ordinary obligations, the fair transaction must lie in the interval between the transactors’ reservation prices.

Weak Non-worseness ensures that possible criteria of fairness can satisfy this constraint. Since it stipulates that at least one Pareto improving transaction must be morally better than not transacting, it also ensures that fair transaction can always be morally better than not transacting. Of course, the original Non-worseness condition can also satisfy this constraint since it makes the even stronger claim that *all* Pareto improving transactions must be morally better than not transacting. However, although welfare is the fundamental concern in exploitation it is not the only one. Our responses to cases of exploitation also imply a strong concern for unfairness. We want to capture the conditional obligation that exploitation presupposes: when more favourable terms are available and are morally obliged, but exploiters fail to transact on these terms, they engage in activity that is morally worse than simply failing to transact. Though the original Non-worseness condition does not allow for such a claim, Weak Non-worseness *does* because it allows some Pareto improving transactions to be morally worse than not transacting, a difference reflected in the shift from ordering (5) to ordering (7). Weak Non-worseness can capture a subtle contextual change that the original condition cannot: it is the presence of the fair, but unchosen option that makes exploiters’ choice to transact unfairly morally worse than not transacting. In other words, they could have done otherwise, but did not.

The intuition that unfair transactions are worse than non-transaction only when a fairer transaction is possible can be captured only by a condition capable of expressing a three place relation between the options. Because Weak Non-worseness expresses a three place relation it can capture the intuition the wrongfulness of exploitation depends crucially upon whether transactors could have done otherwise. Thus, Weak Non-worseness manages to incorporate a concern for fairness, with the fundamental concern for the welfare of the would-be exploited.

---

<sup>33</sup>Robinson (1962: 45)

## 5 Conclusion

I began with two conflicting claims. The first—captured by Pareto Permissibility—is the idea that strict Pareto improvement and consent were jointly sufficient for a transaction to be permissible. The second—captured by Moral Impermissibility—denies that strict Pareto and consent are jointly sufficient for permissible transaction. Necessary Fairness provided a justification for the rejection of Pareto Permissibility; that is, one reason to believe strict Pareto and consent are not sufficient for permissible transaction is that transactions may also be impermissible when they are unfair. I suggested that the primary motivation for the rejected Pareto Permissibility condition was a concern for transactors’ welfare, which can be captured by the Non-worseness condition.

Together, Non-worseness and Necessary Fairness capture a concern for fairness and a concern for welfare, both of which inform our ordinary intuitions about exploitation. Additionally, along with with Betterness of Permissibility, the conditions allow for goodness ordering (5), which explains the debate between those who support sweatshops on welfarist grounds and those who criticise sweatshops for being exploitative.

Unfortunately, these three conditions also *conflict* with the common intuition Pareto improvement does not entail a moral obligation to transact. This intuition is captured by the Optionality of Ordinary Transactions condition. In the previous section I argued that the best way to solve the paradox, while incorporating both a concern for fairness and a concern for welfare, is to replace the original Non-worseness condition with Weak Non-worseness.

## References

- Cohen, G. (1995). *Self-ownership, freedom, and equality*. Cambridge: Cambridge University Press.
- Elster, J. (1982). Roemer versus Roemer: A comment on “New directions in the Marxian theory of exploitation and class”. *Politics and Society*, 11(3): 363–73.
- Goodin, R. (1987). Exploiting a situation and exploiting a person. In *Modern theories of exploitation*. ed. A. Reeve: 166–97. London: Sage.
- Hare, R. (1952). *The language of morals*. Oxford: Clarendon Press.
- Hausman, D. (1995). The impossibility of interpersonal utility comparison. *Mind* 104(415): 473–90.
- Heuer, U. (2011). The paradox of deontology, revisited. In *Oxford studies in normative ethics*. ed. M. Timmons: 236–67. Oxford: Oxford University Press.
- Kristof, N. (2009). Where sweatshops are a dream. *New York Times*, January 14, 2009: 14D.
- Norcross, A. (2006). Reasons without demands: Rethinking rightness. In *Contemporary debates in moral theory*. ed. J. Dreier: 38–54. Oxford: Blackwell.
- Nozick, R. (1974). *Anarchy, state, and utopia*. Oxford: Blackwell.

- Oosterbeek, H., Sloof, R., and van de Kuilen, G. (2004). Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics* 7(2): 171–88.
- Peterson, M. (2013). *The dimensions of consequentialism: Ethics, equality and risk*. Cambridge: Cambridge University Press.
- Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press.
- Robinson, J. (1933). *Economic philosophy*. Bungay: Penguin.
- Roemer, J. (1982a). *A general theory of exploitation and class*. Cambridge, MA: Harvard University Press.
- Ross, R. (2004). *Slaves to fashion: Poverty and abuse in the new sweatshops*. Ann Arbor, MI: University of Michigan Press.
- Sample, R. (2003). *Exploitation: What it is and why it's wrong*. Boulder, CO: Rowman and Littlefield.
- Scheffler, S. (1988). *Consequentialism and its critics*. Oxford: Oxford University Press.
- Sen, A. (1979). Utilitarianism and Welfarism. *The Journal of Philosophy*, 76(9): 463–89.
- Sobel, D. (2007). The impotence of the demandingness objection. *Philosopher's Imprint*, 7(8): 1–17.
- Sobel, D. (2013). Self-Ownership and the conflation problem. In *Oxford Studies in Normative Ethics* Vol 3. ed. M. Timmons: 98–122. Oxford: Oxford University Press.
- Steiner, H. (1984). A liberal theory of exploitation. *Ethics*, 94(2): 225–41.
- Steiner, H. (1987). Exploitation: A liberal theory amended, defended and extended. In *Modern theories of exploitation*. ed. A. Reeve: 132–48. London: Sage.
- Steiner, H. (2010). Exploitation takes time. In *Economic theory and economic thought: Essays in honour of Ian Steedman*. eds. J. Vint et. al.: 20–29. London: Routledge.
- Tversky, A. and Kahneman, D. (1981). “The framing of decisions and the psychology of choice”. *Science* 211(4481): 453–58.
- Veneziani, R. (2013). Exploitation, inequality, and power. *Journal of Theoretical Politics*, 25(4): 526–545.
- Vrousalis, N. (2013). Exploitation, vulnerability, and social domination. *Philosophy & Public Affairs*, 41(2): 131–57.
- Wertheimer, A. (1996). *Exploitation*. Princeton: Princeton University Press.

- Wertheimer, A. (2007). Review: Ruth J. Sample, “Exploitation: What it is and why it’s wrong”. *Utilitas*, 19(2): 259–61.
- Wood, A. (1995). Exploitation. *Social Philosophy and Policy*, 12(2): 135–58.
- Zwolinski, M. (2007). Sweatshops, choice, and exploitation. *Business Ethics Quarterly*, 17(4): 689–727.