Exploitation

Benjamin Ferguson and Hillel Steiner

August 21, 2014 (10160)

(DRAFT: Do not circulate or quote without permission.)

1 Introduction

Broadly speaking, the term 'exploitation' can be applied to a number of different kinds of interaction. Mineral deposits, mistakes, workers, and wombs are all said to be exploited. One sense of exploitation, often applied to the extraction of resources, or capitalising on an opponent's mistaken move in a board game, connotes mere use or advantage taking. Exploitation in this sense need not imply any form of distributive injustice, or indeed, any wrongdoing at all. However, when we claim that a group of workers is exploited we use a second sense of the word. Exploitation in this normative sense involves taking *unfair* advantage and is a form of distributive injustice. In what follows we shall focus solely on exploitation in the normative sense.

When Alice steals Bob's wallet, she takes it without his consent and makes him worse off than he was. Conversely, when Bob buys an apple from Alice's fruit stand, the transaction is consensual and mutually beneficial. Transactions like the apple purchase, in which all parties better off and no parties worse off, are strictly Pareto improving transactions. Some accounts of exploitation include non-consensual and/or harmful transactions such as theft, while others limit their scope to consensual and strictly Pareto improving exchanges. At least regarding normative analysis, it matters little whether we extend the term beyond consensual and mutually advantageous cases, for it is easier to say what has gone wrong when consent is lacking and harm has been done. Explaining the injustice of transactions that are both Pareto improving and consensual, on the other hand, is more difficult.

In fact, some philosophers and economists have claimed that the reason it is difficult to explain the injustice of such transactions is that they are simply not impermissible. They argue,

(1) If a transaction is Pareto improving and consensual, then it is not impermissible.

After all, they reason, if everyone involved is better off transacting and everyone freely agrees to transact, what grounds can there be for condemning the transaction? If a Pareto improving transaction does not occur, then both parties will be worse off than they otherwise would be.

While (1) has some *prima facie* plausibility, it is difficult to endorse in the face of some particular transactions that many people feel *are* impermissible.

Sweatshop labour, for example, is often condemned for being exploitative. Although some sweatshop labour may be harmful to the labourer (and so violate the Pareto condition) or forced (and so be nonconsensual) many sweatshop contracts are both Pareto improving and consensual. Yet it still seems there is something morally amiss with these transactions. Sweatshop labour appears to provide an example of a transaction that is both Pareto improving and consensual, but nevertheless impermissible. That is,

(2) There exist some transactions that are Pareto improving, consensual, and impermissible.

Clearly (2) conflicts with (1). One reason for defending (1) is welfarist. After all, as we noted above, if the transaction does not occur, both parties will be worse off. But what reasons might we have for accepting (2)? What else might be required in order for a transaction to be permissible?

A common response from those who offer accounts of exploitation is that a transaction may also be impermissible if it is unfair. [1] Pareto-superiority and consent are not sufficient for the permissibility of a transaction; transactions must also be *fair* in order to be permissible. While this response sounds reasonable, in order to justify an acceptance of (2) and a rejection of (1) we require an account of what it means for a transaction to be unfair. This is precisely what accounts of exploitation attempt to provide.

1.1 Two Kinds of Fairness

If exploitation is a matter of taking unfair advantage, then Alice exploits Bob only if the benefit she receives from transacting is greater than it ought to be and the benefit Bob receives is less than it ought to be. Though the transaction is mutually beneficial, the distribution of these benefits is unfair. What each transactor's gain 'ought to be' is determined by a particular criterion of fairness. Different accounts of exploitation utilise different fairness criteria to ground exploitation claims.

We can call the feature of the transaction that allows Alice to extract greater benefit from Bob than she would otherwise have extracted 'Bob's disadvantage'. So, suppose Bob has forgotten to eat breakfast and he finds himself in front of Alice's fruit stand, desperate for an apple. Seeing that Bob's demand for an apple is extraordinary, Alice doubles her price. In this case, Bob's hunger is his disadvantage and this disadvantage allows Alice to gain more than she would have from Bob had he not been so hungry. And, of course, in this scenario Bob gains less than he would have because his hunger means he must pay more for the apple. Whether this transaction is an exploitation will depend on whether the difference in the terms of the transaction caused by Bob's hunger is unfair, according to some criterion of fairness.

Existing accounts of the unfairness involved in exploitative transactions take two general forms: they may be *historical* or *ahistorical*. According to historical accounts of unfairness, in order to determine whether the difference in the distribution of the benefits is unfair we need to know what causes the disadvantage. We can partition the historical causes of Bob's disadvantage into three categories. Either Bob was disadvantaged by a person, or he was not. If he was not, then we may say his disadvantage is 'natural disadvantage'. For example, Bob may be naturally disadvantaged because, unforeseeably, a tree fell on his home during a storm. On the other hand, if Bob was disadvantaged by a person, then either it was another person, or it was himself. We may call these forms of disadvantage 'self-caused' and 'other-caused' respectively. So, Bob's disadvantage is other-caused if Alice sets his home on fire and it is self-caused if Bob burns his own home. These forms of disadvantage are shown in the diagram below.

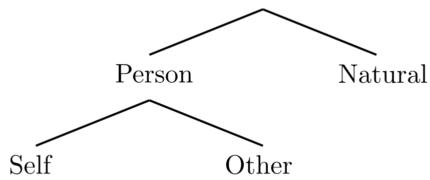

Most—though not all—historical accounts draw on an approach to distributive justice that has become known as 'luck egalitarianism'. For present purposes, we may consider luck egalitarianism in its sloganistic form, according to which "it is bad when persons are worse off than others through no fault of their own." Luck egalitarianism implies that both other-caused disadvantage and natural disadvantage are bad from the point of view of justice. Thus, historical accounts of fairness based on luck egalitarian approaches to distributive justice imply that if Bob's disadvantage is either other-caused disadvantage or natural disadvantage then it is, broadly speaking, unfair when this disadvantage allows Alice to gain more—and Bob less—than either would have gained absent the disadvantage. So, if Alice gains more from Bob because Carol has *stolen* Bob's breakfast, then Alice's higher price is unfair because Bob's disadvantage was other-caused. The same holds if Alice's gain is brought about by Bob's natural disadvantage; perhaps his breakfast blew away in the wind. However, luck egalitarian approaches to distributive justice also imply that if Alice is able to charge Bob more for the apple because he forgot to eat breakfast, then her higher price is *not* unfair. This is because Bob's disadvantage was self-caused. To reiterate, for historical accounts of fairness, the source of disadvantage is relevant for determining the fairness of the transaction. Often, but not always, historical accounts rely on broadly luck egalitarian theories of distributive justice.

According to ahistorical accounts of unfairness, the cause of the disadvantage is either always irrelevant (which is a strong form of ahistoricism), or irrelevant in certain contexts (which is a weaker form) for determining whether the thereby modified terms of a transaction are unfair. One defender of the strong form of ahistoricism is Robert Goodin, who claims that exploitation consists in a failure to protect the vulnerable (i.e. disadvantaged) in our interactions with them "regardless of the particular source of their vulnerability" (Goodin 1987, p. 187). Alternatively, Ruth Sample defends a weaker version of ahistoricism, claiming that we exploit others when we fail to consider whether their basic needs are met when we transact with them. For Sample this obligation is not predicated on the reasons these persons' basic needs are unmet and therefore, is ahistorical. However, unlike Goodin, her focus is limited to disadvantage that falls below a certain threshold. When persons' basic needs are unmet, *regardless of the reason* they are unmet, we ought to constrain our advantage over them, according to Sample. Thus, her criterion of fairness implies a commitment to a weak form of ahistoricism. According to ahistorical accounts of fairness, the unfairness of exploitation is not—or not only—located in prior distributive injustice. Rather, it is grounded in the (non-historical) properties of the *present* transaction.

1.2 Two Kinds of Advantage Taking

Because their focus is on the features of the present transaction, ahistorical accounts often place an emphasis on the attitude, intentions, or motive of the exploiter. Transactions ordinarily described as exploitations often seem to involve a blatant disregard for the exploited's situation or a disrespect for their person. This advantage-taking feature should be distinguished from the maldistribution—or unfairness—involved in exploitation. Approaches to exploitation differ not only according to whether they employ historical or ahistorical criteria of fairness, but also according to their interpretations of 'advantage taking'.

One interpretation of what it means to take advantage, endorsed by Jon Elster, is rather minimalist. Elster claims we take advantage of another whenever we "[add] an option to his opportunity set that is better both for him and for [us] than the option preferred by him prior to the expansion" (Elster 1982, p. 364). On this account Alice can be said to 'take advantage' of Bob (and vice versa) simply by engaging in a mutually beneficial transaction with him.[2] Accounts that employ Elster's minimal sense of advantage taking will characterise exploitation simply as an unfair, but Pareto improving transaction. What makes the transaction *unfair* will, of course, depend upon the notion of fairness the account endorses.

On the other end of the spectrum lie accounts that claim advantage taking involves engaging in such transactions with attitudes, intentions, or motivations that are morally amiss. For example, it may be claimed that in order for Alice to exploit Bob the transaction in which they engage must be not only unfair, but she must be aware that it is so and yet fail to constrain her advantage over Bob.[3]

These two distinctions—whether exploitation's unfairness is to be understood ahistorically or historically and whether the presence of inappropriate attitudes is necessary for exploitation—can, in theory, be variously combined. However, in practice, accounts of exploitation that employ ahistorical criteria of fairness generally also place a strong emphasis on the attitudes of the exploiter towards the exploited and accounts that employ historical criteria of fairness tend to favour Elster's account of advantage taking.

	Attitudes Unnecessary	Attitudes Necessary
Historical Fairness	Steiner Roemer	Ferguson
Ahistorical Fairness	Marx	Goodin Sample

In the sections that follow, the distinction between historical and ahistorical accounts of fairness takes centre stage. Section two focuses primarily on Robert Goodin's (1987) and Ruth Sample's (2003) ahistorical approaches to exploitation. Both accounts also place a strong emphasis on the attitudes involved in advantage taking. The section will also briefly consider the classical Marxist account of exploitation. Though Marx is not concerned with the attitudes of exploiters, his account of exploitation is based upon an account of fair transaction that can be construed as ahistorical. Section three addresses two historical approaches proposed by John Roemer (1982a; 1982b) and Hillel Steiner (1984; 1987; 2010; 2013). Each of these approaches implicitly endorses Elster's minimalist account of advantage taking. Section three concludes with a brief outline of an historical approach proposed by Benjamin Ferguson (2013) that also includes an attitudinal approach to advantage taking. The final section revisits the tension between the claim in (1) that Pareto improvement and consent are sufficient for the permissibility of a transaction and the claim in (2) that some transactions that are both Pareto improving and consensual are, nevertheless, impermissible.

2 Ahistorical Accounts of Fairness

As we noted in the introduction, exploitative transactions are unfair transactions and this unfairness may, broadly speaking, be understood in two ways. Transactions can be unfair in an historical sense when Alice does better than she would have—and Bob worse than he would have—because of a prior unjust distribution of advantage. Alternatively, transactions can be unfair in an ahistorical sense when they diverge from a criterion of fairness that does not make reference to the justice of Alice and Bob's distributive shares prior to entering the transaction. Rather, the transaction is judged fair or unfair according to its own (non-historical) properties or to the non-historical properties of the transacting parties situations. Let's revisit Alice and Bob's apple transaction. Suppose that the transaction is unfair. This means a fair transaction is one in which Bob pays less than he does in the actual transaction. There are many ahistorical criteria of fairness that we might use justify the claim that the terms of the actual transaction are unfair to Bob. For example, we might believe that a fair transaction is one in which both transactors receive equal utility gain from transacting. Alternatively, Marx argued that Alice exploits Bob when the amount of labour that went into making the apple is less than the amount of labour it took Bob to earn the wages he used to purchase it. Equal utility gain and equality of embodied labour time are both ahistorical fairness criteria because each refers to the value of the goods exchanged in the present transaction. In fact, they are very similar ahistorical criteria since each argues that the exchange is fair only when it is *equal.* They differ according to the currency they use as the equalisandum.

Alternative ahistorical criteria of fairness do not require equality in the exchange. Rather, they imply that a transaction is fair when it surpasses a particular threshold. According to these criteria, Alice's transaction with Bob is unfair when—on Sample's account—she fails to take Bob's basic needs into account, or—as Goodin argues—when she fails to act in a way that protects Bob from his vulnerability to her in the transaction. In what follows we first consider

Goodin's strong form of ahistoricism, followed by Sample's more constrained approach.

2.1 Goodin's Account

According to Goodin's account, exploitation in general involves "playing for advantage in situations where it is inappropriate to do so . . . [it] consists in playing games of strategy in circumstances which render them somehow inappropriate" (Goodin 1987, p. 184). He identifies four kinds of situation where playing for advantage is inappropriate: when playing against those who have "renounced playing for advantage themselves", or "who are no match for you in games of advantage", or "who are unfit or otherwise unable to play in games of advantage at all", or "when your relative advantage derives from others' grave misfortunes" (Goodin 1987, p. 184–187). All four situations are, he claims, "manifestations of one particular kind of wrong," namely, a failure to honour the moral norm of "protecting the vulnerable. . . *regardless of the particular source* of their vulnerability" (Goodin 1987, p. 187, emphasis added). Thus, for Goodin exploitation occurs when we use others' vulnerabilities to inappropriately secure some form of benefit in games of strategy, irrespective of the vulnerabilities' source. Goodin's claim that any form of vulnerability can generate exploitation places his account clearly in the ahistorical camp.

Goodin's understanding of vulnerability is central to his account of exploitation. He claims vulnerability is both relational—"a full specification will tell us who is vulnerable to whom with respect to what"—and relative—"A is more vulnerable to B (1) the more control B has over outcomes that affect A's interests and (2) the more heavily A's interests are at stake in the outcomes that B controls" (Goodin 1986, p. 118). That is, A is vulnerable to B with respect to x when B has the ability to control x and x has some impact on A.

While this analysis is promising as an account of vulnerability, it threatens to make the conception of exploitation that is based upon it too wide in scope. The claim that we ought to constrain our advantage when we interact with vulnerable others is convincing in many cases. Surely Alice is wrong to use her advantage to extract labour from Bob for very low wages in a sweatshop. But is she also obliged to constrain her advantage over Bob if he is a robber baron who, after making his millions, has "renounced playing for advantage"?[4] What if her advantage over Bob derives from a grave misfortune that Bob himself is responsible for bringing about? If Bob is vulnerable to Alice just in case Alice can control some outcome or action that impacts Bob, it is unconvincing to claim that all cases in which Alice fails to make allowance for this vulnerability amount to exploitation. After all, we are willing to spend more money on food when we are hungry, more willing to purchase products if we have seen them advertised, and more desperate for a new job if we are currently unemployed. In the broadest sense, each of these situations creates vulnerability, yet it is not clear that it is inappropriate to take advantage of all of these vulnerabilities. Indeed, all Pareto improving transactions involve vulnerability in Goodin's sense. If Alice owns an apple that Bob wants, then regardless of the reason Bob wants the apple, it is true that he is vulnerable to her with respect to the apple since she has the ability to control the apple and whether or not Bob's obtaining the apple has some impact on him. Yet, far from being impermissible, in most cases, taking advantage of Pareto improving transactions should be encouraged. Of course, *some* of these transactions may be unfair, but this unfairness must consist in something other than taking advantage of vulnerability in Goodin's sense. It seems that Goodin fails to explain what is special about those transactions we ordinarily call exploitative.

Nevertheless, Goodin is correct that vulnerability or disadvantage is *necessary* for exploitation. If Bob were not in some way vulnerable or disadvantaged vis-à-vis Alice, then Alice could not gain more and he less than either ought to gain from the transaction. However, as the previous paragraph makes clear, some form of vulnerability is present in nearly every transaction, yet not every transaction is an exploitation. Thus, though Goodin identifies a necessary condition for exploitation, it is not *sufficient*. Perhaps, in light of this potential overextension, we should search for a notion of vulnerability that is more constrained in scope. What we are after is a *restricted* sense of vulnerability of which it is more intuitively 'unfair' to take advantage. For such an approach, we turn to Ruth Sample's account of exploitation.

2.2 Sample's Account

Sample argues that exploitation in general involves "interacting with another being for the sake of advantage in a way that degrades or fails to respect the inherent value in that being" (Sample 2003, p. 57). She claims these failures of respect fall into three domains: we can fail to show proper respect for other persons by neglecting what is necessary for the person's basic needs, by taking advantage of past injustices, or by engaging in a transaction involving "an aspect of that person's being that ought not to be commodified" (Sample 2003, p. 57).

According to Sample's first domain of disrespect, we ought to constrain our advantage over others when those with whom we transact are so impoverished that their basic needs are unmet. We treat others unfairly when we fail to respond appropriately to their impoverishment and this is so regardless of the reason these needs are unmet. Since the historical reasons for others' impoverishment play no role in generating this obligation, this domain employs an ahistorical notion of fairness. However, unlike Goodin's unconstrained concept of vulnerability, Sample's appeal to a threshold level of wellbeing is more limited in scope. That we ought to constrain our advantage in the face of severe poverty does not necessarily imply that we ought to constrain our advantage when we transact with those who are less poor, but nevertheless vulnerable in Goodin's wider sense.

Sample also argues that we ought to constrain our advantage even when interacting with the vulnerable who are not impoverished. However, this obligation is limited to situations in which their vulnerability is the result of past injustice. Sample's second domain of disrespect employs a *historical* notion of fairness. Whether we ought to constrain our advantage when interacting with those whose basic needs are met, but are nevertheless vulnerable, does depend upon the *reason* for their vulnerability. These domains provide two distinct reasons for Alice to constrain her advantage. For example, suppose a very hungry Bob approaches Alice's stand to purchase an apple. If Bob is starving and his basic needs are unmet, then according to Sample's first domain, Alice ought to constrain her advantage over Bob—perhaps by offering him the price she would extend to a non-hungry customer.[5] On the other hand, if Bob's basic needs are met (but he is still quite hungry), then the *reason* he is hungry becomes important. Though Sample does not endorse any particular theory of distributive injustice, we can consider the example in light of a luck egalitarian account. In this case, if Bob is hungry because Carol stole his lunch, then Alice ought to constrain her advantage, but if Bob is hungry because he forgot to eat breakfast, then Alice is not obliged to constrain her advantage. Of course, these two domains may also be combined: Bob's basic needs may be unmet because of a past injustice. Taking advantage of persons in these situations seems, on Sample's account, to be particularly unjust.

The final domain of disrespect in Sample's account is disrespect that issues from trading in non-commodifiable goods, such as human organs or sexual intercourse. While it is true that such transactions are occasionally referred to as exploitative, it is less clear that these kinds of transaction are exploitations in the sense we have been considering. Without denying that they may be degrading or disrespectful, it does not seem that transactions falling *only* within this domain are cases of unfair advantage taking. Indeed, Sample acknowledges that when the other forms of disrespect are not present, but a person violates the non-commodification condition they are "not exploited... however, [they] may be degraded if the appropriate case against commodification can be made" (Sample 2003, p. 83). For this reason, in the remaining discussion we will restrict our attention to Sample's first two forms of disrespect.

There are many points in favour of Sample's account. It captures the intuition that we ought to constrain our advantage in light of others' vulnerabilities, but it avoids the overextension of Goodin's account by placing a restriction on this obligation when the vulnerabilities are less severe, that is, in cases where our interactors' basic needs are met. Nevertheless, there is one concern that arises not only in the case of Sample's account, but also for Goodin's account. The concern is that inasmuch as these ahistorical (or partly ahistorical) accounts imply unconditional duties on the part of the would-be exploiter not to take advantage of his or her superior position in certain situations, these duties are *themselves* exploitable.

2.3 The Moral Hazard Objection

Suppose Bob is a wealthy banker who takes a calculated gamble. If the gamble goes well, Bob will gain a great deal, but if it does not, his basic needs will go unmet. According to Sample's first form of disrespect, if Bob loses the gamble and his basic needs are unmet, then Alice is obliged to constrain her advantage when transacting with him. Failure to do so amounts to exploitation according Sample's conception. But suppose Alice's obligation is known to Bob prior to his decision to take the gamble. In particular, he knows that if the gamble goes well, the large gain is his to keep, but if it goes poorly, his loss will be subsidised by Alice since she is obliged to constrain her advantage over him.[6] The obligation for Alice to constrain her advantage in the event of Bob's loss creates a moral hazard that Bob can use to reduce the cost of losing the gamble. In effect, Alice's obligatory subsidy is a form of insurance for Bob's gamble that he does not pay for. Alice's obligation is tantamount to a moral vulnerability that, if enforced, would also be an empirical vulnerability. Taking advantage of the vulnerability created by the moral hazard is unfair and—since exploitation involves taking unfair advantage—it seems appropriate to say that Bob exploits Alice when he does so.[7]

According to the basic needs domain, Alice exploits Bob if she fails to constrain her advantage while transacting when his basic needs are unmet. Yet, according to the above argument, when Bob takes advantage of the moral hazard by claiming this advantage from Alice, he exploits her. This situation creates a conflict between two mutually exclusive and exhaustive acts.[8] A desideratum of any normative theory is that it provides noncontradictory guides to action. Yet here the theory implies that the subsidy is both obligatory and impermissible. It cannot be the case that both conflicting obligations are valid; either Alice is not obliged to constrain her advantage or Bob is not obliged to constrain his advantage (which is bestowed upon him by her having such an obligation).[9]

The problem with Sample's claim that Alice must constrain her advantage over Bob—even when Bob is responsible for his own vulnerability—is that it is insensitive to personal responsibility. The *reason* Bob is vulnerable is relevant to determining whether Alice is obliged to constrain her advantage. Requiring Alice to constrain herself when Bob is responsible for his own vulnerability makes *Alice* vulnerable to being exploited by Bob. A solution to this tension is to deny that Alice exploits Bob when she fails to constrain the advantage she gains from Bob's self-caused vulnerability. Indeed, concerns about this form of exploitation provide part of the motivation for luck egalitarians' omission of self-caused disadvantage from the domain of injustice.

Of course, this luck egalitarian claim is controversial. Critics of the approach have advanced the 'unacceptable outcomes objection', which claims "the luck egalitarian embrace of responsibility implausibly justifies leaving people without assistance in thoroughly bad situations if they are responsible for bringing them about" (Stemplowska 2009, p. 251).[10] Though we believe the objection is misplaced, Goodin disagrees, arguing "[o]nce all their opportunities for self-help have passed... the situation is beyond their control. Others, however, may still be able to act so as to avert harm to them. To suggest that those others should (or even that they *may*) stand idly by and watch people reap the bitter fruits of their own improvidence is surely absurd" (Goodin 1986, p. 129).

It is important to note though that we can both accept the unacceptable outcomes objection—agreeing with Goodin that we are obliged to aid those whose disadvantage is self-caused—*and* reject the claim that Alice *exploits* Bob if she fails to constrain her advantage when his vulnerability is self-caused. We may claim that Alice is obliged to constrain her advantage *despite the fact* that when she subsidises Bob she allows him to treat her unfairly. We do not need to abandon the intuition that what Bob does is unfair in order to claim that Alice is nevertheless obliged to provide aid. Rather, it is possible to claim that a concern for Bob's welfare, now that his basic needs are unmet, overrides a concern for fairness. What we claim here is that it is incorrect to use the concept of exploitation to explain the obligation Alice has towards Bob because, as the above example shows, exploitation is a responsibility-sensitive concept. If these arguments are correct, then Goodin's and Sample's omission of personal responsibility in their accounts is misplaced. Whether Alice exploits Bob will depend, importantly, on the historical *reasons* Bob is disadvantaged.

In addition to the moral hazard objection a second, practical, concern arises for both Goodin's and Sample's accounts. Though Goodin tells us we ought to protect all those who are vulnerable and Sample claims that exploitation involves a failure to respect those with whom we interact, neither provides a concrete account of how the obligations to protect or respect others should be operationalised. Recall Alice's sale of an apple to Bob. Let us suppose that the sale price of \$2 for the apple is unfair because it violates Sample's second domain of disrespect. By selling Bob the apple for \$2, Alice takes advantage of Bob's vulnerability that was caused by past injustice. Though Sample's domain explains the wrongfulness of Alice's action, it does not specify what price would ensure Alice's interaction with Bob was respectful. Should she sell the apple for \$1? For \$1.50? Of course, the 'fair' price will likely depend upon some contextual details that have not been specified in this case, but a complete account of exploitation should at least provide some general principles that can be applied to specific contexts. This concern is not an objection to either of the two accounts, rather, it represents an important detail that must be specified if the accounts are to properly guide our decisions in transactions.

2.4 Marx's Account

One account that does provide a concrete guide to when a transaction is fair is the classical Marxist account of exploitation. According Marx, workers are exploited when they work more hours than are needed to produce the goods they consume. So, if Bob works for Alice for 10 hours, but the wages he is paid by Alice only allow him to purchase a bundle of goods that comprise 8 hours of work, Bob is exploited by Alice. Thus, according to Marx a transaction is non-exploitative just in case the labour time embodied in the goods exchanged is equal.

Since Marx's concern was primarily with employment contracts, the scope of his account is somewhat narrower than the accounts we have considered so far. Nevertheless, the idea that people should receive remuneration that is equivalent to the amount of work they do is normatively compelling. Marx's account of fairness, like those of Goodin and Sample, can be seen as an ahistorical account. The transaction is unfair when the labour time embodied in the exchanged items is unequal *regardless* of the historical reasons for this inequality. However, Marx's account does not suffer from the same moral hazard problem that plagues Goodin's and Sample's accounts because the amount of embodied labour Alice is obliged to transfer to Bob depends on the amount of labour Bob provides for Alice. The reliance on equality of labour exchanged ensures that there is no distributive unfairness in the transfer between Alice and Bob. Unfortunately, the account suffers from other problems. In particular, the labour theory of value that underpins the equal exchange notion of fairness in Marx's account has long been criticised for being difficult to operationalise.

Determining whether the amount of embodied labour transferred between Bob and Alice is equal requires a measure of labour. However, a convincing measure of embodied labour has been notoriously difficult to come by. Clearly a fair exchange is not one in which the 'raw' amount of time laboured is equal, for some workers may be more talented than others. In order to incorporate differences in labour, Marx argues that the equalisandum is not mere 'raw' labour, but 'socially necessary labour'. Yet this move simply passes the definitional buck to 'socially necessary labour'. We now need a criterion for distinguishing socially necessary labour from unnecessary labour. This lack of a convincing definition of socially necessary labour, along with some other well known problems for Marx's labour theory of value—its inability to reliably explain the relationship between labour value and prices and the arbitrariness of labour as the value numéraire—have led most economists to reject Marx's labour theory of value.[11]

Two critical responses to the classical Marxist account of exploitation have emerged in light of these difficulties with the labour theory of value. The first is espoused by Robert Nozick who concluded that 'with the crumbling of the labor theory of value, the underpinning of [the Marxist] theory of exploitation dissolves" (Nozick 1974, p. 253). For Nozick, the labour theory of value was essential to the Marxist conception of exploitation. The second critical response originates with G. A. Cohen. Cohen argues that though the labour theory of value is untenable, "what really animates Marxists, whatever they may claim…is not that the capitalist gets some of the value the worker produces [as the labour-theory implies], but that he gets some of the value *of what* the worker produces…[workers] do not create *value*, but they create *what has value*" (Cohen 1979, p. 354). Nozick noted that the particular form of exploitation found in the classical account is formulated in terms of the labour theory of value, and that this theory of value is false. But Cohen points out that amended versions may succeed where Marx's original account fails. Cohen's observation heralded a research programme that sought to provide alternative accounts of equal exchange based on equalisanda other than equality of embodied labour time, many of which also employed historical accounts of fairness. In the following section we consider two such accounts, John Roemer's equality-based account of exploitation and Hillel Steiner's rights-based account.

3 Historical Accounts of Fairness

The two accounts of exploitation that we focus on in this section–John Roemer's property relations account and Hillel Steiner's rights-based account—draw upon Marx's idea that exploitation involves *unequal* exchange, but they diverge from the classical Marxist account by offering different accounts of what must be equalised in order for a transaction to be fair. Unlike Sample and Goodin these two accounts employ historical approaches to fairness. Both accounts also adopt Elster's minimal notion of advantage taking, though, as we will see at the end of the section, this adoption poses problems for both accounts.

3.1 Roemer's Account

Roemer begins his property relations account by noting that we can think of a group as being exploited if it has some hypothetical "conditionally feasible alternative under which its members would be better off", which he captures in the following two conditions (Roemer 1982b, p. 276). A coalition S in a larger society N is exploited if and only if:

(R1) There is an alternative, which we may conceive of as hypothetically feasible, in which S would be better off than in its present situation.

(R2) Under this alternative, the complement to S, the coalition $N - S = S'$, would be worse off than at present (Roemer 1982a, p. 195).

Various forms of exploitation can be specified in Roemer's model by altering the withdrawal conditions that specify the hypothetical situations. Roemer focuses on three withdrawal conditions that, he claims, capture feudal, capitalist, and socialist exploitation. Under specific withdrawal conditions coalitions are exploited if (R1) by withdrawing from society and creating their own economy they can achieve utility levels that dominate the utility frontier available to the coalition when it remains in the original economy *and* (R2) the coalition's complement fares worse in this respect if the coalition withdraws (Roemer 1982a, p. 196). Roemer outlines the following three withdrawal conditions:

> **Feudal exploitation.** The coalition can take with it its own endowments of both alienable (transferable, nonhuman property) and inalienable (i.e., skills and talents) resources.
>
> **Capitalist exploitation.** The coalition can take with it its *per capita* share of society's alienable property.
>
> **Socialist exploitation.** The coalition can take with it its *per capita* share of all endowments, alienable and inalienable (Roemer 1982a).

For example, taking sweatshop labourers as coalition S, we can say that they are capitalistically exploited according to Roemer's model if and only if the utility they could achieve by withdrawing from society and creating their own economy *with their per capita share of assets*, is greater than the utility they gain by working in sweatshops *and* the utility of others in society (S')—including the sweatshop owners—is less after the labourers withdraw.

An attractive feature of Roemer's account is that it separates the structure of an exploitative transaction (conditions R1 and R2) from the ethical content that allows us to condemn the transaction (the withdrawal conditions). The normative withdrawal conditions of exploitation are specified by various approaches to distributive justice, while the structural conditions specify the way in which these distributive elements are manifest in an exchange. Roemer's framework allows us to agree about whether coalitions are capitalistically or feudally exploited even if we disagree about whether the withdrawal conditions that characterise these forms of exploitation are normatively compelling.

However, this does not mean Roemer's account entirely escapes criticism. Though few philosophers have expressed doubts about the *necessity* of Roemer's conditions, many argue that R1 and R2 are not *sufficient* to establish the existence of an exploitative transaction. While exploitation does require Alice to gain and Bob to lose, this gain and loss must come about *in the right way.*

One concern is that R1 and R2 do not ensure that an exploiter's gain comes at the expense of the exploited. Imagine two islands societies, S and S', each unknown to the other. Suppose S has \$5 of assets and S' has \$15. If the assets of the islands were combined ($\$5 + \$15 = \$20$) and the inhabitants of each were allowed to withdraw with their per capital share of assets ($\frac{\$20}{2} = \$10$), (R1) and (R2) would imply that S' exploits S since S would be better off with \$10 and S' would be worse off with \$10. Yet, it does not seem appropriate to say that S *exploits* S'. This is because the wealth of S' does not depend on—come at the expense of—S action. Roemer recognises this problem and adds a further condition: "capitalist exploitation can be said to exist only if one coalition is gaining at the expense of another" (Roemer 1982c, p. 300) that is, "if the would-be exploited withdrew with *their own* endowments, the complementary coalition would be worse off" (Sensat 1984, p. 28).

Though this condition brings Roemer's account closer to capturing the relation of exploitation between the two groups, Jon Elster argues that the account is still insufficient. Elster asks us to consider two groups that have ideological interactions, but no economic interaction. One group, S is hedonistic and prefers leisure to work. The other, S', has a puritanical work ethic and disdains the hedonism of S. This disdain causes S' to work harder than they otherwise would, leading to a greater productivity for S'. The hedonists in S, on the other hand, find the puritanical work ethic contemptible and this attitude encourages their hedonism. Consequently, the presence of S' negatively affects the productivity of S. If S' did not exist, then S would be better off. If S did not exist, then S' would be worse off. Conditions (R1) and (R2) are satisfied. But it seems absurd to claim that S' exploits S precisely because S' does not appropriately cause S to be worse off (Elster 1982, p. 368). Roemer acknowledges that certain "bizarre" counterexamples may be produced unless the following 'dominance' condition is added:

(R3) S' is in a relationship of dominance to S.

However, as Kymlicka, Bertram, and even Roemer himself point out, the addition of R3 is *ad hoc*. Roemer acknowledges that "since dominance is undefined, and is as elusive a concept as exploitation, the addition of (R3) is *ad hoc*" (Roemer 1982b, p. 277, fn. 15). Even if an analysis of dominance were provided, the condition would remain *ad hoc*—at least for Roemer's purposes—for two reasons. First, as Kymlicka points out, appeals to dominance or power are "disconnected from the 'ethical imperative' [Roemer] identifies as the basis of exploitation theory", namely, distributive injustice (Kymlicka 1990, p. 204, fn. 13). Second, as Bertram notes, "the different withdrawal conditions were supposed to bring to light the normative differences which exist between Marxists and neo-classicals. The invocation of a further condition is not going to satisfy someone who considers that the forced transfer of income through taxation to finance welfare programmes *is* exploitative" (Bertram 1988, p. 126).

So, although the first problem of dependence can be handled by requiring that one coalition gain *at the expense* of the other, the problem that a lack of dominance presents for Roemer's account is more substantial. But, we hasten to point out that it is not fatal. The presence of R3 is indeed *ad hoc* if our goal is, like Roemer's, to provide a purely distributive account of exploitation. However, it may turn out that exploitation is not, in fact, a purely distributive concept. An adequate conception of exploitation may need to move beyond Elster's minimalist account of advantage taking to include a richer analysis that makes room for concepts like power, domination and intention. A modified version of Roemer's account that offers an analysis of what it means for S' to be in a relationship of dominance to S may successfully narrow the scope of Roemer's initial two conditions. Such an account would move from the top-left quadrant of the grid we presented in section 1.2 to the top-right quadrant.

Finally, we would like to mention a pragmatic reason for moving away from Roemer's account. Roemer's work on exploitation includes not only the property relations account we have outlined here, but also an 'equal exchange' account, which, we have suppressed for reasons of space.[12] Together both of Roemer's accounts have generated a large number of interesting results for Marxist theory. However, Roemer's focus on hypothetical economies and the exploitation of *coalitions* S and S' makes it difficult to determine whether a single transaction between Alice and Bob is exploitative. Roemerian exploitation is an aggregate phenomenon that is not easily applied to individual transactions. If we wish to explain how Marx's concern with class struggle can be analysed in terms of exploitation and inequality without relying upon the labour theory of value, then Roemer's account is extremely useful. However, if our aim is to explain how individual transactions can be Pareto improving, consensual, but unfair, another account may be more useful. With this in mind we turn now to Hillel Steiner's rights-based account of exploitation.

3.2 Steiner's Account

Steiner's guiding aim is to show how exploitation can be understood as a form, specifically, of *injustice* (Steiner 1984, 1987, 1994, 2010, 2013). Despite its long-standing association with the idea of injustice, he observes, the currently numerous theories of distributive justice are all devoid of any account of exploitation. To remedy that lacuna and exhibit that connection, Steiner begins by focusing on the paradigmatic form of injustice—violations of moral rights.

He asks us to imagine an auction where Bob is selling some objects or services, Xs, and Alice enters the winning bid of \$1.50.[13] Under what conditions is it plausible to regard those Xs as having been undersold? That is, under what conditions can we say that an X's \$1.50 price is unjustly low and *should* have been higher? A common response, especially prior to the advent of neoclassical economics in the 1870s, has been to claim that an X is worth more than \$1.50 because the cost of producing it was greater than \$1.50. Although this objective conception of value is endorsed by both classical Marxist and recent cost-of-production (Reiff 2013) accounts, Steiner shares modern mainstream economics' rejection of this response, inasmuch as it relies upon the pre-neoclassical view that things exchanged have a value which is not determined by their sellers' and buyers' preferences. That is, he accepts persons' *actual* preference orderings as the determinants of prices. Those preferences are taken as given: they are not to be gainsaid by reliance upon some notion of value that is independent of them. And, moreover, they cannot be gainsaid by any liberal theory of justice that aspires to be neutral as between different conceptions of 'the good life'.

This does not, however, oblige us to accept an X's \$1.50 price as just. It does not do so because, under different circumstances, it might have sold for a higher price, say, \$2.00. So the question is: What are those counterfactual circumstances, and how does their absence make the \$1.50 price unjust?

One such circumstance, we might be tempted to think, is that others—someone else or Alice herself—have preferences that place a higher value than \$1.50 on an X. But this counterfactualised circumstance is insufficient to prove that \$1.50 is an unjust price. Justice does not require anyone to value an X at all, let alone at more than \$1.50. As previously indicated, we are to take persons' actual preferences as given.

Suppose there is someone, Carol, who *does* value an X at more than \$1.50. Why, then, did she not outbid Alice? One reason might be that she couldn't afford to do so. But, again, although this explains why Alice won, it doesn't suffice to render her \$1.50 price unjust. What we need to know is why Carol couldn't afford a higher bid. Why was her wealth endowment insufficient for that? Suppose her endowment was insufficient because she had previously lost some of it in a risky investment venture. That, however, would not sustain a claim that the \$1.50 price was unjust. It was *Carol's* actual preference to undertake that risk and, hence, the potential loss it entailed. She committed no injustice in doing so and her incurring that loss was not unjust.

But suppose, instead, that the reason why Carol's endowment was insufficient is that, on her way to the auction, she was robbed. That is, had she not been robbed—had she not been unjustly deprived of some of what she had a moral right to possess—her outbidding Alice would have occurred and Bob's Xs would have sold for, say, \$2.00 each. It is by virtue of that prior injustice—that unjust reduction of Carol's endowment—that Alice's \$1.50 bid was successful and thereby deprived Bob of the additional \$.50 per X that he would have acquired in the absence of that injustice. Bob, in short, is *exploited* to the tune of \$.50 per X.

This, then, represents the core of a conception of exploitation that can be clearly and directly construed as it commonly is construed: namely, as a form of injustice. An exploitation is a transaction in which the exploitee (Bob) gains less—and the exploiter (Alice) more—than they would each have gained had a prior injustice not occurred. What can count as such a prior injustice will, of course, depend on some wider theory of distributive justice. Each individual's entitlements—what they have just rights to possess and what they are justly permitted to do with those possessions—will vary according to which particular account of distributive justice (Rawlsian, Dworkinian, left-libertarian, Nozickian, Rothbardian, etc.) one endorses. We can amplify this conception of exploitation by indicating a further range of variables that might count as instances of that prior injustice.

So, for instance, an alternative reason why Carol did not outbid Alice might be that she was forcibly prevented from attending the auction. *If*, according to some theory of distributive justice, Carol's moral rights include various claims, liberties, powers and immunities that entitle her to participate in the auction, her forcible prevention from doing so amounts to an injustice. Similarly, if her failure to attend the auction was due to someone depriving her of a notice of the auction's occurrence—a notice which she was morally entitled to receive—that, too, could amount to an injustice.

More complexly, if someone imposed a surcharge on the selling price of an X, that *could* amount to an injustice. Suppose there was an imposed surcharge of 10%. Then, although Carol was prepared to pay up to \$2.00 for each of Bob's Xs, she was unwilling to pay that extra 10%, i.e. \$2.20. Alice's winning bid of \$1.50 incurs the same surcharge, and she has to pay \$1.65. In this case, Bob gets only \$1.50 for each of his Xs; the surcharger gets \$.15; and Alice gets Bob's Xs whereas, in the absence of that surcharge, she would have had to bid more than \$1.50 to acquire them. What is happening here is that the surcharger is, in effect, exercising partial ownership over the wealth endowments of the bidders.[14] Whether we should consider that exercise to be an injustice—that is, whether the bidders are justly entitled to full ownership of their endowments—depends, of course, on the particular theory of distributive justice we endorse.

Yet another form of prior injustice, sufficient to entail that the \$1.50 price is exploitative, consists in what may previously have happened to the exploitee, Bob, himself. So far, the prior injustices we've considered have all been ones committed against a third party, Carol. It is the violation of Carol's rights—with regard to her wealth endowment and her auction participation—that has supplied the condition necessary for Bob's exploitation by Alice. But, equally, it can be a prior violation of the moral rights of Bob himself that supplies that condition. Thus, suppose that, prior to the auction, Bob himself was robbed.

Before that robbery, Bob's *reservation price* for his Xs was, say, \$1.80 each: he would not have sold the Xs for less than \$1.80 each.[15] The effect of the robbery, in diminishing his endowment, is to lower his budget line to the point where his reservation price is only, say, \$1.20. It is thus by virtue of that robbery that Alice is able, exploitatively, to purchase his Xs for as little as \$1.50 each.

Steiner's conception of exploitation as injustice thereby analyses it as consisting in a pair of sequenced events: a prior injustice which consists in the violation of a relevant person's rights—the exploitee's or some third party's—followed by the exploitative transaction itself.[16]

3.3 Advantage Taking

Both Steiner's account and Roemer's accounts of exploitation are based upon historical accounts of fairness: for both, a prior injustice is necessary for a subsequent exploitation. The two accounts also share a certain amount of flexibility with respect to the theories of distributive justice that underpin this prior injustice. Just as Roemer's account is compatible with a variety of withdrawal conditions, so too does Steiner's account allow for a variety of theories of distributive justice. Indeed, if the theory of distributive justice that is employed in Steiner's account is the same as that which gives content to the withdrawal conditions in Roemer's account, then the scope of the resulting conceptions of exploitation will be very similar.

Suppose that the withdrawal conditions for Roemer's account are resource egalitarian: agents may withdraw with their per capita share of alienable endowments. Then a coalition S is exploited if, after withdrawing, Roemer's conditions R1 and R2 are satisfied. That is, if the members of the coalition had their per capita share of alienable assets they would be better off and their compliment, S', would be worse off. Similarly, if a resource egalitarian theory of distributive justice is combined with Steiner's account, then Alice exploits Bob in a transaction just in case Bob gains less, and Alice more, than they each would have gained because of a prior unequal distribution of alienable resources. The scope of Steiner's and Roemer's conceptions differ (rather trivially) because Roemer's conditions focus on coalitions, while Steiner's are expressed in terms of individuals and (more significantly) because of the evaluations of whether the person or coalition is better or worse off. Roemer's evaluation of whether a coalition is better (worse) off depends on whether, in the post-withdrawal economy, there is a utility frontier available to the coalition that dominates (is dominated by) the frontier available prior to withdrawal. Reaching this frontier may require a number of transactions between the agents in the coalition. On the other hand, on Steiner's account, a person is better (worse) off if they gain more (less) utility from an individual transaction, rather than a series of transactions. Roemer's account focuses on how a coalition fares in the long run, making it more appropriate for questions related to the state of particular socioeconomic classes while Steiner's focus on how agents fare in particular transactions makes his account more appropriate for evaluating the state of individuals.

So, while the extensions of the two historical accounts are not identical, they are very similar. This similarity suggests a possible problem for Steiner's account. If the extension of Roemer's account, sans a dominance condition, is too wide, then it seems the same could be true of Steiner's approach. And indeed, such a criticism has been raised by Steven Walt, who argues "...if a title to objects resulting from a rights violation is invalid and exploitation involves a rights violation, then it follows that most if not all unequal bilateral exchanges are cases of exploitation. But this is implausible. The presence of exploitation is unquestionable. Its omnipresence is questionable" (Walt 1984, p. 243). The problem Walt identifies is common to many historical accounts of justice. A past rights violation may infect all 'downstream' transactions, making the holdings resulting from all subsequent transactions unjust. Since there is good reason to believe that there exists at least one rights violation in the history of transactions leading up to Alice's sale of an apple to Bob, there is good reason, on Steiner's account, to believe not only that this transaction is exploitative, but indeed, that most transactions are exploitative. In this case, the extension of Steinerian exploitation exceeds our ordinary understanding of exploitation.

One reason we may think the scope of exploitation characterised by Steiner's and Roemer's accounts outstrips our ordinary understanding of exploitation is that, though the accounts capture an unfair—or unjust transaction—the conditions that the accounts provide do not ensure that this unfairness comes about "in the right way"; that is, in a way that justifies our calling the transactions exploitative. Both accounts seem to allow for a kind of 'accidental exploitation'. However, we may think that exploitation should be tied more tightly to the attitudes or intentions of the exploiter. What exercises us about sweatshops, for example, is not merely that some individuals happen to gain unfairly at the expense of others, but rather that they seek out vulnerable individuals from whom they can be certain to extract large profits. Though Roemer and Steiner both provide compelling analyses of the *unfairness* involved in exploitation, their approaches do not seem to account for the sense in which exploiters *take advantage* of their victims.

In fact, in this sense it seems that the ahistorical accounts fare much better. Though Goodin's claim that any form of vulnerability could generate obligations to constrain our advantage was too broad, his remark that exploitation involves playing for advantage in an inappropriate way seems to be on the right track.

One approach, suggested by Ferguson (2013), is that a conception of exploitation that aligns more closely with our ordinary understanding of the concept can be provided by adding an 'awareness' condition to historical approaches like Steiner's and Roemer's.[17] On this approach exploitation differs from mere unfair transaction because exploitation involves cases in which Alice both gains unfairly at Bob's expense and has a true belief—is 'aware'—that she does so. When Alice sees that Bob is unjustly disadvantaged and yet, in light of this awareness fails to constrain her advantage over him, not only does she gain unfairly from Bob, but she actively *takes advantage* of him. Advantage taking in this sense, then, is stronger than Elster's minimalist claim that we take advantage of others when we engage in a mutually beneficial transaction with them. Not only does the 'awareness condition' limit the scope of these theories, but Ferguson's account can be seen as a hybrid account that incorporates concerns from both ahistorical and historical approaches. It combines the thicker notion of advantage taking implied by ahistorical accounts' concern for the attitudes of the exploiters towards those they exploit with the concern for personal responsibility incorporated in the historical accounts' approach to fairness.

As a hybrid account, however, Ferguson's approach has its drawbacks. In particular, the satisfaction of the awareness condition depends on persons' beliefs; thus, on Ferguson's account, whether an act is exploitative depends on whether it has a certain *intensional* description. However, as Steiner has argued, "whether an act is just *doesn't* depend on... its having a certain intensional description. It depends, rather, on its having a certain extensional description: that is, on whether it's compatible with a certain distribution of freedom, with a certain set of rights" (Steiner 1994, p. 212). In this sense, the link between exploitation and injustice is weakened. On Ferguson's account exploitation is not a specific form of injustice, rather, it becomes a *wrongful*, or impermissible response to prior injustice.

4 Conclusion

As was noted in the Introduction, some philosophers and economists believe that

(1) If a transaction is Pareto improving and consensual, then it is not impermissible.

Contrarily, it is widely claimed that

(2) There exist some transactions that are Pareto improving, consensual, and impermissible.

All the theories of exploitation examined here support (2) and reject (1): that is, all exploitative transactions are morally impermissible. Whether we should regard *legal interference* with such transactions as morally warranted will thus depend on (i) how closely our conception of exploitation is associated with that specific form of moral impermissibility that we call *injustice*, and (ii) whether we believe that injustices are the only form of morally impermissible acts that warrant legal interference.

References

Bertram, Christopher. 1988. "A Critique of John Roemer's General Theory of Exploitation." *Political Studies* 36: 123–130.

Bronfenbrenner, Martin. 1973. "Samuelson, Marx, and Their Latest Critics." *Journal of Economic Literature* 11: 58–63.

Cohen, Gerald. 1979. "The Labour Theory of Value and the Concept of Exploitation." *Philosophy and Public Affairs* 8: 338–360.

Eatwell, John, Peter Newman, and Murray Milgate, eds. 1987. *The New Palgrave: A Dictdionary of Economics*. London & New York: Macmillan.

Elster, Jon. 1982. "Roemer Versus Roemer: A Comment on *New Directions in the Marxian Theory of Exploitation and Class*." *Politics and Society* 11: 363–373.

Ferguson, Benjamin. 2013. "The Paradox of Exploitation: A New Solution." Ph.D. thesis, London: LSE.

———. 2014. *Exploitation and Disadvantage.* Unpublished Draft.

Goodin, Robert. 1986. *Protecting the Vulnerable.* Chicago: University Of Chicago Press.

———. 1987. "Exploiting a Situation and Exploiting a Person." In Reeve 1987: 166–200. London: Sage.

Guest, Anthony. 1961. *Oxford Essays in Jurisprudence*, First Series. Oxford: Oxford University Press.

Honoré, Anthony. 1961. "Ownership." In Guest 1961: 107–147.

Kymlicka, Will. 1990. *Contemporary Political Philosophy: An Introduction.* Oxford: Clarendon Press.

Nozick, Robert. 1974. *Anarchy, State, and Utopia.* Blackwell: Oxford.

Reeve, Andrew, ed. 1987. *Modern Theories of Exploitation.* London: Sage

Reiff, Mark. 2013. *Exploitation and Economic Justice in the Liberal Capitalist State.* Oxford: Oxford University Press.

Roemer, John. 1982a. *A General Theory of Exploitation and Class.* Cambridge, MA: Harvard University Press.

———. 1982b. "New Directions in the Marxian Theory of Exploitation and Class." *Politics and Society* 11: 253–287.

———. 1982c. "Property Relations vs. Surplus Value in Marxian Exploitation." *Philosophy & Public Affairs* 11: 281–313.

Sample, Ruth. 2003. *Exploitation: What It Is and Why It's Wrong.* Boulder: Rowman and Littlefield.

Samuelson, Paul. 1971. "Understanding the Marxian Notion of Exploitation: A Summary of the So-Called Transformation Problem Between Marxian Values and Competitive Prices." *Journal of Economic Literature* 9: 399–431.

Sensat, Julius. 1984. "Exploitation." *Nous* 18: 21–38.

Steedman, Ian. 1987. *"Reservation Price and Reservation Demand."* In Eatwell et al. 1987: vol. 4, 158–159.

Steiner, Hillel. 1984. "A Liberal Theory of Exploitation." *Ethics* 94: 225–241.

———. 1987. "Exploitation: A Liberal Theory Amended, Defended and Extended." In Reeve 1987: 132-148.

———. 1994. *An Essay on Rights.* Oxford: Blackwell.

———. 2010. "Exploitation Takes Time." In Vint, et al. 2010: 20–29.

———. 2013. "Liberalism, Neutrality and Exploitation." *Politics, Philosophy & Economics* 12: 335–344.

Stemplowska, Zophia. 2009. "Making Justice Sensitive to Responsibility." *Political Studies* 57: 237–259.

Vint, John, J. Stanley Metcalfe, Heinz D. Kurz, Neri Salvadori, and Paul Samuelson, eds. 2010. *Economic Theory and Economic Thought: Essays in Honour of Ian Steedman.* London: Routledge.

Walt, Steven. 1984. "Comment on Steiner's Liberal Theory of Exploitation." *Ethics* 94: 242–247.

Wertheimer, Alan. 1996. *Exploitation.* Princeton: Princeton University Press.

Notes

[1]Different accounts of exploitation offer different accounts of *why* the transaction is unfair. According to the classical Marxist account of exploitation, the transaction is unfair if the socially necessary labour time embodied in the goods being transferred is unequal. According to Steiner's account, an unfair distribution of advantage leads to an unfair distribution of gains from transaction. According to Alan Wertheimer's account (1996), transactions are unfair when their terms differ from those that would be offered in hypothetical competative markets. We discuss these and other accounts of fairness in the sections that follow.

[2]Note that advantage taking in Elster's sense is not morally amiss. Rather, advantage taking in this sense becomes exploitative when the Pareto improving transaction is also unfair.

[3]See Ferguson (2013) for an outline and defence of an account that takes this approach.

[4]Thanks to Mike Otsuka for drawing our attention to this point.

[5]What exactly showing proper respect for Bob requires, in terms of transaction prices, is left unspecified in Sample's account.

[6]Alice's constraint of advantage can be characterised as a subsidy since that proportion of the gains from transacting she refrains from claiming accrues to Bob.

[7]Note that Alice's obligation to Bob is a *conditional* obligation. Alice is not obliged to transact with Bob, but *if* she does, then she is obliged to constrain her advantage over him. However, provided the transaction is (strictly) Pareto improving, Alice will *prefer* to transact with Bob. So, though Alice is not morally obliged to transact, she is prudentially advised to do so, and if she does, then she must constrain her advantage over Bob. It is this conditional obligation that suffers from the moral hazard objection.

[8]The two possible acts are Alice's provision of the subsidy and Bob's refusal of the subsidy when the two transact. If Alice provides the subsidy, then Bob has not refrained from accepting it; if Bob refrains from accepting the subsidy, then Alice has not provided it. The two acts are mutually exclusive. Since Alice's failure to provide the subsidy ($\neg\phi$) is equivalent to Bob's refusal to accept the subsidy ($\neg\neg\phi$, i.e., ϕ) the two acts are exhaustive.

[9]The moral hazard problem is explored in more detail in Ferguson (2014).

[10]It should be noted that Stemplowska favours the luck egalitarian account and argues *against* the objection.

[11]For a discussion of the arbitrariness of labour as the value numéraire see Roemer (1982a). The transformation problem is discussed in Samuelson (1971) and an overview of the debate between Samuelson and Marxist economists is

provided in Bronfenbrenner (1973).

[12]A detailed treatment of Roemer's unequal exchange account can be found in Roemer (1982a). Roemer also provides a very useful (and less formal) summary of this work in Roemer (1982b). Elster offers a critique of Roemer's property relations approach and endorses the unequal exchange account in Elster (1982).

[13]Auctions are markets writ small.

[14]On the composite nature of ownership, and its consequent interpersonal divisibility, see Honore (1961).

[15]Cf. Steedman (1987) 158: "The simplest example of a reservation price is that price below which an owner will refuse to sell a particular object in an auction. Since the owner could always, in principle, enforce such a price by outbidding everyone else, this leads immediately to the more general concept of a reservation price as that price at which the owner of a fixed stock will choose to retain some given amount from that stock, rather than supply more. . . . ".

[16]Since some of those prior injustices constitute unjust reductions in persons' endowments, it follows that an exploitative transaction can also be the result of prior exploitative transactions. If Bob receives less than he is justly entitled to for his Xs, the bid that he can enter at a subsequent auction, for Doug's Ys, will be less than it would have been had he (Bob) not been exploited in his sale of the Xs.

[17]Indeed, the awareness condition can be seen as an operationalisation of Roemer's dominance condition.