

Are we all exploiters?*

Benjamin Ferguson^{1,2}

¹VU Amsterdam

²University of Warwick

Correspondence

Benjamin Ferguson, Department of
Philosophy, University of Warwick and
Department of Philosophy, VU Amsterdam
Email: benjamin.ferguson@warwick.ac.uk

Abstract

This paper argues that two single-factor accounts of exploitation are inadequate and instead defends a two-factor account. Purely distributive accounts of exploitation, which equate exploitation with unfair transaction, make exploitation pervasive and cannot deliver the intuition that exploiters are blameworthy. Recent, non-distributive alternatives, which make unfairness unnecessary for exploitation, largely avoid these problems, but their arguments for the non-necessity of unfairness are unconvincing. This paper defends a two factor account according to which A exploits B iff A gains unfairly from B and either A believes that the gains he receives in the transaction wrong B, or A is culpably unaware that the gains he receives in the transaction wrong B. This account avoids the problems of non-distributive approaches and also delivers the intuition that exploiters are blameworthy.

1 | INTRODUCTION

Recall the last transaction that you completed. Perhaps you bought breakfast on your way to work, or paid for public transportation, or purchased an item of clothing. Surprisingly, according to nearly all moralised accounts of exploitation, that transaction was wrongfully exploitative. These accounts entail that either you wrongfully exploited the seller or the seller exploited you (but it is hard to say

* This paper was presented at VU Amsterdam and the Frankfurt School of Finance and Management. I thank those present for their suggestions. I also thank Chris Thompson, Sebastian Köhler, Chris Ranalli, Hillel Steiner, Christine Tiefensee, Marina Uzunova, Roberto Veneziani, Alex Voorhoeve, Jan-Willem Wieland, and an anonymous referee and editor at *Philosophy and Phenomenological Research* for helpful comments.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2020 The Authors. *Philosophy and Phenomenological Research* published by Wiley Periodicals LLC on behalf of Philosophy and Phenomenological Research Inc

which). If this claim is correct, then nearly all accounts of exploitation make the concept so pervasive that it appears to become explanatorily and morally empty.

Now, exploitation is not necessarily wrong. As Allen Wood (1995, 2016) points out, it is not wrong to exploit your opponent's move in competitive games. Some accounts of exploitation are non-moralised and simply offer a descriptive account of the kinds of social relations that must hold between agents for exploitation to obtain. For example, many commentators argue Marx's claim that "workers are exploited if they work longer hours than the number of labour hours embodied in the goods they consume" (Elster 1986: 121) is a descriptive, non-moralised account.¹ Although non-moralised accounts do not entail that exploitation is necessarily wrong, they can, of course, claim that it is wrong in cases where certain wrong-making features are present.

The most prominent accounts of what makes exploitation wrong, when it is wrong, appeal to unfairness. These accounts begin with the claim that wrongful exploitation involves 'taking *unfair* advantage' (Arneson 2013, 2016; Elster 1982; Ferguson 2018; Ferguson and Steiner 2018; Goodin 1987; Reiff 2013; Sample 2003; Steiner 2010; Wertheimer 1996; Zwolinski 2012; Zwolinski and Wertheimer 2016) and ground the wrongfulness of exploitation in unfairness. They then proceed to offer an account of what it means for a transaction to be unfair. I will call these analyses *purely distributive* accounts because they appeal *solely* to a maldistribution of goods that obtains between agents. In other words, purely distributive accounts are those that accept the following claim:

Equivalence. Wrongful exploitation is tantamount to unfair transaction.

The equivalence claim does not, on its own, make wrongful exploitation pervasive. Wrongful exploitation is pervasive because the most plausible accounts of transactional fairness entail:

Pervasiveness. Unfair transactions are pervasive.

If equivalence and pervasiveness are both true, then wrongful exploitations are pervasive. But this seems odd. If nearly all transactions are unfair, then what is the point of singling out particular practices, such as commercial surrogacy (Wertheimer 1996; Wilkinson 2003), sweatshops (Arnold and Bowie 2003), and price gouging (Zwolinski 2008) as wrongfully exploitative? After all, *your* last transaction was likely wrongfully exploitative too. As Nicholas Vrousalis (2016: 537) puts it, when " 'unfair treatment' and 'exploitation' are used interchangeably what is the extra purchase of saying that A exploits B?" There is something *morally distinctive* about these paradigmatic cases of exploitation and the possibility that wrongful exploitation is pervasive conflicts with the ordinary intuition that it applies to a limited range of cases. Call this the *scope problem*.

Now, you might not be too troubled by the scope problem. It is plausible that despite common intuitions, some forms of wrongdoing *are* widespread. For example, Peter Singer (1972) has forcefully argued we don't do enough for the poor, even though saving many lives is easy for those who are wealthy. The conclusion that the wealthy act wrongly does not seem to be diminished by the pervasiveness of the wrongdoing.² However, in the case of exploitation, if the pervasiveness and equivalence claims are true, then not only will nearly *every transaction* be a wrongful exploitation, but as I will argue, in many cases we will not be able to identify which of the parties gains unfairly in a

¹Allen Wood (1972; 1981) is perhaps the most well known proponent of the claim that Marx's is a non-moralised account. Norman Geras (1985) offers a helpful summary of the debate.

²I thank an anonymous referee for suggesting this example.

transaction, despite having good reason to believe the transaction is unfair. So, although we can blame the wealthy in Singer's case because we have a pretty good idea of who the wealthy are and who the poor are, the way in which unfairness is pervasive not only makes it difficult to distinguish exploitative and non-exploitative transactions, it also makes it difficult to distinguish exploiters from the exploited parties. Therefore, it is often inappropriate to assign blame in many cases of unfair transaction. And if the equivalence thesis is true, and wrongful exploitations are just unfair transactions, then the same holds for wrongful exploitation. The pervasiveness of wrongful exploitation also conflicts with the ordinary intuition that exploiters are morally responsible and blameworthy for exploiting. Call this the *blame problem*.

I will argue that paradigmatic cases of wrongful exploitation³ involve more than the *actus reus* wrong of unfairness. They also involve a second, *mens rea* wrong: the exploiter's failure to respond to a moral reason to constrain his gain. I call this two factor account the *epistemic conception* of exploitation because in addition to unfairness it also requires that exploiters are, or ought to be, occurrently aware that the gains they receive in the transaction wrong the other party. Unlike purely distributive accounts, the epistemic conception can avoid the scope and blame problems.

In the next section I argue the pervasiveness claim about unfairness is not only true, but true in a way that makes it hard to identify who gains an unfair share. Avoiding the scope and blame problems requires rejecting the equivalence claim and thus, purely distributive accounts of exploitation. In section three I consider and reject two ways that paradigmatic cases of exploitation might differ in *degree* from merely unfair transactions before critiquing Nicholas Vrousalis's (2013) account of how they could differ in *kind*. In section four I present my two factor account, according to which exploitation involves *both* unfair transaction and an appeal to the epistemic condition for moral responsibility.

2 | THE PERVASIVENESS CLAIM

Purely distributive accounts are motivated by a variety of antecedent moral and political positions. But they share a deflationary understanding of the 'advantage taking' component and, thus implicitly (and in many cases, explicitly) accept the equivalence claim. I suspect that many people find the equivalence claim relatively uncontroversial, but they are suspicious of the pervasiveness claim. They doubt that unfair transactions are really all that pervasive. If you are sceptical, let me try to convince you.

The most prevalent approaches to transactional fairness are substantive accounts (Zwolinski and Wertheimer 2016; Ferguson and Steiner 2018), which claim the unfairness of a transaction is a feature of *what* is agreed to. These contrast with procedural accounts that identify transactional unfairness with a procedural defect in the transaction.

Most substantive accounts equate fair transaction with a particular distribution of the social surplus created by transacting.⁴ Some claim a fair transaction is one in which the gains are equally distributed (Ferguson and Ostmann 2018; Ostmann 2017), others appeal to rational bargaining solutions (Nash

³Hereafter, unless otherwise stated, the scope of my analysis will refer to *wrongful* exploitation.

⁴Basic needs accounts (Sample 2003; Snyder 2008), which claim a transaction is unfair when it involves a failure to ensure the basic needs of a transactor are met, are an exception. However, these cannot be accounts of unfairness since meeting basic needs is both insufficient for fair transaction (since there is nothing conceptually incoherent about claiming that a transaction between two relatively wealthy persons is unfair) and unnecessary for fair transaction (since it is implausible that transactions between two poor persons, neither of whom can meet the other's basic needs are unfair *by definition*).

1950; Kalai and Smorodinsky 1975). Now, recall your last transaction. Were the gains in that transaction precisely equally divided or divided in conformity with a rational bargaining solution? Likely they were not. And, *if* they were not, then by these surplus-based accounts of fair transaction and the equivalence claim, the transaction was an exploitation. Furthermore, note that even though these accounts give us reason to believe the transaction was an exploitation, in most cases we do not know who exploited whom because we do not know who gained the unfair share. And this means that in many cases we cannot ascribe moral responsibility for exploitations.

Perhaps you are not convinced by my appeal to *your* last transaction, so let me offer one of mine as another example. Most mornings I commute by bus. Let's say a return ticket costs \$5. Now, although my willingness to pay for this ticket varies—I may have a class to teach, the weather might be bad, or good—the price of the ticket does not change. Yet if my reservation price shifts each day, then *however* it is determined, the price that represents a fair distribution of the gains should also shift. Since the price does not shift, most of these transactions are unfair (sometimes to me, sometimes to the bus company). And if they are unfair and the equivalence claim is true, they are wrongful exploitations.

You might respond that the problem lies not with the equivalence claim, but rather with substantive accounts of fair transaction. After all, since substantive accounts only focus on the outcome distribution that results from transactions and not on the background conditions and context in which transactions occur, they are insensitive to important determinants of fairness such as personal choices, merit, and historical injustice. This means they not only conflict with intuitions about desert, but they also give rise to moral hazard problems (Ferguson 2016).

Let us instead consider procedural accounts. Alan Wertheimer argues that transactions are fair when they reflect the prices that would obtain in hypothetical, moderately idealised markets. The fair price is “the price that an informed and unpressured seller would receive from an informed and unpressured buyer” (Wertheimer 1996: 230). Wertheimer argues that when goods exchange at these prices, no party “takes special unfair advantage of particular defects in the other party's decision-making capacity or special vulnerabilities in the other party's situation” (Wertheimer 1996: 232). However, it is not clear that fairness requires the elimination of informational asymmetries and pressure, since information “is itself a commodity that has value, and arguably, those who invest in gathering information are entitled to a return on their efforts” and not all pressure is inappropriate: “A may be willing to pay B more for an apple if he forgot to eat breakfast... But fairness does not require that B sell A an apple for the price that A would have been willing to pay had he eaten breakfast” (Ferguson 2018: 9). Perhaps the most charitable way to interpret the core of Wertheimer's idea is that the fair price is the price that would obtain in a *just* market. After all, terms that result from *unjust* pressure, informational asymmetries, and vulnerabilities are surely unfair.

On this reading, Wertheimer's account is remarkably similar to Hillel Steiner's (1984; 1987) and John Roemer's (1982a; 1982b) accounts of exploitation. According to both, A exploits B when a prior unjust endowment of resources makes A better off and B worse off in a transaction than either would have been had the prior distribution of assets been just.⁵ An advantage of procedural approaches is that

⁵For Steiner, A and B are individuals, for Roemer they are coalitions of individuals. Roemerian exploitation is only defined for groups (though it can be reinterpreted for individuals). This means that *no* individual transactions will be exploitations and thus, when applied to groups, Roemerian exploitation is not pervasive. However, it still encounters the blame and scope problems. Consider scope. We cannot call individual labour contracts exploitations if exploitation is not defined over individual transactions. Sweatshops workers *are* members of an exploited group, but so too are rag pickers (and all other poor persons). Thus, we cannot explain what's *distinctive* about sweatshops. Now take blame. We cannot blame *individual transactors* for exploitation if exploitation is not defined over individual transactions. And we can only blame an exploiting group if that group acts as a group agent. But plausible conditions for group agency (List and Pettit 2011) are often unmet in cases where we'd nevertheless like to ascribe blame.

they make fairness depend on a prior theory of distributive justice (Bajaj 2015; Ferguson and Steiner 2018). Unlike substantive accounts and Wertheimer's account, their appeal to a prior theory of justice makes them sensitive to a variety of factors relevant to fairness, such as individual effort and prior wrongs.

Nevertheless, procedural accounts also entail that unfair transaction is pervasive. Indeed, I am not the first to point out this consequence for Steiner's account. As Steven Walt (1984: 242) put it,

[M]ost if not all titles have a far from 'clean' causal ancestry ... hence 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.

Although Walt's quote appeals to historical entitlement theories, the same point can be made for procedural accounts based on alternative theories of justice. If unjust advantages are endemic, unfair transaction is pervasive. There is good reason to think unjust advantages *are* endemic. Therefore, we should accept the pervasiveness claim. Most transactions, including those that are rather pedestrian, are very likely unfair to one of the parties involved. So, if you think that paradigmatic instances of wrongful exploitation somehow differ in kind from these more commonplace forms of unfair transaction, it *appears* you must reject the equivalence claim.

3 | THREE UNTENABLE SOLUTIONS

In much of this section I focus on Nicholas Vrousalis's account of exploitation as domination, which offers an explanation of how wrongful exploitations differ in *kind* from unfair transactions. But first I briefly consider two ways they might differ only in degree. Perhaps paradigmatic cases involve greater degrees of unfairness or they involve particularly vulnerable persons (or both) and so they stand out as more egregious instances of exploitation than 'ordinary' unfair transactions.

First, consider unfairness. While those who work in sweatshops are paid unfair wages and are almost certainly the victims of injustice, it is not obvious that the differential between what they would have earned against just background conditions and what they currently earn is *greater* than those in their societies who are similarly poor and work in other industries. Or, to appeal to surplus-based accounts, it is not clear that the *proportional* maldistribution of gains in cases of, say, price gouging, is greater than the maldistribution present in other daily transactions.

Second, consider vulnerability. Sweatshop workers are very vulnerable. But sweatshops generally offer better conditions and pay than occupations such as rag-picking that are otherwise available to these persons (Ferguson 2016; Powell and Zwolinski 2011). If most transactions are unfair and wrongful exploitation is tantamount to unfair transaction, then rag-pickers are also commonly exploited. Yet, even though they are more vulnerable than sweatshop workers, the way in which rag-pickers are exploited seems to differ from the way that sweatshop workers are exploited. So the degree of a person's vulnerability also does not seem to distinguish the paradigmatic cases. In general these degree-based solutions fail to address what Jonathan Wolff describes as the "one nagging doubt... that, somehow, there just seems to be more to exploitation than unequal exchange" (1999: 107). So let us now focus on how exploitations might differ in *kind* from unfair transactions.

Nicholas Vrousalis's domination account is the only prevalent account of wrongful exploitation that explicitly rejects the equivalence claim, and, consequently, avoids the scope and blame problems.

Vrousalis writes, “on my view, exploitation is not tantamount to unfair advantage-taking, for there can be fair exploitings” (Vrousalis 2013: 149n). The truth of the pervasiveness claim provides a reason to think unfair transaction is not sufficient for exploitation, but Vrousalis goes further, arguing that it is not necessary. According to Vrousalis’s account, “A economically exploits B if and only if A and B are embedded in a systematic relationship in which A instrumentalizes B’s economic vulnerability to appropriate (the fruits of) B’s labor” (2013: 138). He later clarifies that exploitation occurs “only when A instrumentalizes B’s vulnerability to enrich herself, *such that A subordinates B*” (Vrousalis 2016: 530). For Vrousalis, exploitation is, at root, related to the quality of the interaction between A and B. Unfair distribution is not essential to exploitation even though it *can* arise as the result of dominating interactions.

This focus on the quality of interaction between the exploiter and exploited is a much more plausible distinction between the pedestrian forms of unfairness that are implied by standard accounts, and the kind of moral concerns that paradigmatic cases of exploitation raise. It is not just that unjust background conditions allow A to gain unfairly at B’s expense, it is that A gains from B by dominating her.

Vrousalis produces a compelling argument against the necessity of unfairness, citing the following three inconsistent claims, “(i) exploitation is unfair advantage-taking, (ii) fairness⁶ is responsibility-constrained equality, (iii) exploitation can arise from any material inequality” (2013: 149, n.54). Vrousalis argues that since “(ii) and (iii) are very compelling... (i) must be rejected” (2013: 149, n.54). That is:

P1	Fairness is responsibility-constrained equality.
P2	Exploitation can arise from <i>any</i> material inequality.
C	Therefore, unfairness is not necessary for exploitation.

Vrousalis supports his argument with the following case:

Rescuer. A finds B in a pit. A can get B out at little cost or difficulty. A offers to get B out, but only if B agrees to pay a million euros or to sign a sweatshop contract with A. B signs the contract. (Vrousalis 2013: 148)

In *RESCUER*, intuitively, A exploits B and, crucially, this intuition does not depend on whether B is responsible for her plight. Since we do not know whether B is responsible, but we do think she is exploited, it seems that P2 must be true. Vrousalis simply stipulates that P1 is the most plausible definition of fairness in a transactional context (and the aforementioned problems for substantive accounts lend support to P1). The argument is valid, so if both premises are true, then unfairness is indeed unnecessary for exploitation (and the equivalence claim is false).

I think most people will agree that, *prima facie*, *RESCUER* is exploitative. However, I also suspect many form the intuition that A exploits B on the basis of their belief that he charges her an *unfair* price. That is, they think any plausible fairness criterion should entail that one million euros for rescue is unfair, and thus, *because it is unfair*, *RESCUER* is exploitation. Of course, if the intuition that *RESCUER* is exploitation stems from a belief that the transaction is unfair, the case doesn’t support the conclusion that unfairness is unnecessary for exploitation. Consider the following variation:

Fair Rescuer. A finds B in a pit. A can get B out at little cost or difficulty. A offers to get B out, but only if B agrees to pay a fair price. B agrees.

⁶There is a typo in the original text. Vrousalis writes ‘unfairness’, but surely he means *fairness*.

The sole difference between FAIR RESCUER and RESCUER is that the price is stipulated to be fair. But since fairness is not necessary for exploitation according to Vrousalis, the variation should not change our intuition that the case is exploitation. Indeed, FAIR RESCUER, not RESCUER provides the crucial test of P2. Because it is hard to see how FAIR RESCUER could be exploitation, it seems we should conclude that unfairness *is* necessary for exploitation.

I say ‘seems’ because Vrousalis’s P1 provides a reason to think that, at least for responsibility sensitive theories of fair transaction, the standard version of RESCUER is entailed by FAIR RESCUER. Assume B is uncontroversially responsible for her plight. Suppose A carefully sealed and fenced off the pit, posted clear warnings about the price he would charge for rescue, and that despite these warnings, B hopped the fence, pried off the cover, and intentionally jumped into the pit. If these conditions (and any others you think are necessary for B’s responsibility) are met, then friends of responsibility sensitive theories of distributive fairness could argue that since B is responsible, A’s price is fair. And thus, they could argue that RESCUER is covered by the scope of ‘fair price’ in FAIR RESCUER.

What are we to make of this result? There are three possible responses. First, if you think it is a mistake to conclude that B’s responsibility entails A’s price is fair, then you can reject the association of fairness with responsibility-constrained equality. That is, you can reject P1. Note that you can even do so without abandoning responsibility sensitivity. Serena Olsaretti (2009) has argued that although the grounds of responsibility concern “the features of people that we can... hold them responsible for”, there is a further question about the fairness of ‘stakes’: the costs “attach[ed] to whatever features constitute the justifiable grounds of responsibility” (Olsaretti 2009: 170). If you take this approach, then you can agree B’s responsibility places some burden on *her* without concluding that all *laissez faire* stakes are fair. Alternatively, if you think P1 is true, then you could bite the bullet and deny P2, accepting that instances of RESCUER in which B satisfies plausible conditions of responsibility are not exploitative. Finally, if you find neither of these options attractive, then you must accept Vrousalis’s conclusion that there can be fair exploitations. But this means Vrousalis’s argument is compelling only for those who adopt responsibility-sensitive accounts of fairness *and* laissez-faire approaches to stakes that are compatible with P1 *and yet* are also unwilling to bite the bullet and deny RESCUER is exploitation. However, it is unlikely that many would accept *both* premises, since doing so would require accepting something like a luck egalitarian or libertarian account of transactional fairness, but repudiating its implications in cases like RESCUER. The argument against the necessity of unfairness is not compelling.

And so, with no strong argument against the necessity of fairness in an account of exploitation, I turn my attention to a two-factor account that includes both unfairness and an additional condition that limits exploitation’s scope.

4 | The TWO-FACTOR ACCOUNT

Solving the scope and blame problems does not require abandoning unfairness. Appending a second condition to unfairness suffices. I argue that what distinguishes paradigmatic cases of wrongful exploitation from mere unfair transaction is the satisfaction of the epistemic condition for moral responsibility (Rudy-Hiller 2018; Sliwa 2017; Robichaud and Wieland 2017). Unlike the purely distributive conception, which appeals only to unfair transaction, this epistemic conception implies that exploiters are morally responsible for exploitation. And, in order to be morally responsible, they must either be *aware* of what they are doing (Wieland 2017), or if they are ignorant, they must be culpable for their ignorance (i.e. negligent). First consider awareness.

In the context of unfair transaction, when can we say that the exploiter is aware? Perhaps the most obvious possibility is that A is aware when he intends to gain unfairly from B. In this case, in place of the equivalence claim, we could say A engages in the epistemic conception of exploitation of B iff A gains unfairly from B *and* A *intends* to gain unfairly from B. An appeal to intention restricts the scope of the purely distributive approach, which addresses the pervasiveness problem and allows us to claim A is morally responsible.

However, many exploitation theorists have argued intentionality is too strong (Sample 2003; Arneson 2016; Steiner 2018). The best objection to an intentionality condition arises from the possibility of side-effect cases. A could engage in a transaction with B from which he *knew* he would gain unfairly, without intending to gain unfairly:

Regional Presence. A opens a retail outlet in a region because a business presence is strategically important. He foresees that economic conditions allow him to pay the shop's few employees an unfair wage, but the strategic geographic considerations are much more important than saving salary expense.

A's *intention* is to establish a presence in the country. The fact that he can also save on salary expense is merely a foreseeable side effect. Yet, surely if A foresees that he will pay the workers unfair wages, he is still morally responsible for this act. It is appropriate to blame A for underpaying his workers in this case, even if he does not *intend* to exploit them. Therefore, intention is too strong.

Suppose instead that A exploits B under the epistemic conception iff A gains unfairly from B and A *knows* he gains unfairly from B. Knowledge is (debatably) justified true belief, plus conditions designed to avoid Gettier's (1963) and related (Pritchard 2012) problems. It is also, like intention, sufficient to both limit exploitation's scope and to capture blameworthiness. However, a knowledge condition faces its own problems. Suppose A *believes* his transaction with B is unfair, it is *true* that the transaction with B is unfair, but A's belief is not justified:

Unfair Tuesdays. A believes it is unfair to sell products at prices ending in odd numbers on Tuesdays. He sells B an apple for \$151 on a Tuesday.

In most scenarios \$151 is an unfair price for a single apple. Let's suppose it *is* unfair by whatever happens to be the correct account of fair transaction. But, of course, the transaction also just so happens to be unfair according to A's own *implausible* theory. Since A forms his belief that the transaction is unfair on the basis of an implausible theory, he is not justified in believing the transaction is unfair. So, despite the fact that the transaction is unfair and that A believes it is unfair, A does not *know* it is unfair. And if he does not know it is unfair, then according to the knowledge-based condition, he is not blameworthy for exploiting A. Knowledge is also too strong.⁷

Weakening the condition further to appeal to belief also encounters problems.

Domination Theorist. Suppose A believes exploitation does not involve unfairness, but only domination *tout court*. A believes he dominates B, he believes transacting with her will exploit her, but as a domination theorist, he does not believe the transaction will be

⁷See Rosen (208) for a similar case. Peels (2014: 493–94) discusses the non-necessity of justification.

unfair. Suppose that A nevertheless transacts with B and his belief that it is a fair transaction is false.

It is hard to deny A engages in blameworthy exploitation. Yet, he does not believe that the transaction is *unfair* and so doesn't seem to satisfy a belief-based condition. Belief also seems too strong. Now we are running out of rope, for it is unclear how belief can be relaxed without altogether abandoning some kind of epistemic condition and falling back on the equivalence claim.

The general structure of the epistemic conditions I have considered is: A has [mental state] that [proposition]. The mental states of intention, knowledge, and belief are all susceptible to counterexamples. However, I have assumed that the content of the proposition variable in the condition is: 'the transaction is unfair'. We saw that A can be blameworthy for exploitation despite not knowing, or intending. However, the domination theorist case does not show that A has *no* beliefs about his action. Rather, it shows that he has the wrong sort of beliefs. That is, the beliefs he has do not satisfy a belief based condition if we assume the propositional variable is completed by 'the transaction is unfair'. Perhaps the counterexample can be handled by focusing on the content of the propositional variable, rather than the mental state variable.

In the domination theorist counterexample A does not believe the proposition 'the transaction is unfair'. But he does believe that the transaction wrongs B. Furthermore, this belief is true. It's just that the belief is true in virtue of the fact that the transaction is unfair and not (as A mistakenly thinks) because he dominates B. The epistemic conception need not require that A believes his gains are *unfair*, *per se*. Rather, he must simply believe that the gains he receives *wrong* B. Thus, we may say that A exploits B under the epistemic conception iff (1) A gains unfairly from B and (2_w) A believes that the gains he receives in the transaction wrong B.

Condition (2_w) requires that A is occurrently aware both of the action he performs (Mele 2010; Sliwa 2017) and of its moral valiance (Zimmerman 2002, 2017; Sliwa 2017). It links A's belief that the transaction is wrong to A's gains. Even though A can be blameworthy for exploitation without believing that the transaction is unfair, exploitation requires that his belief that the transaction is wrong is linked to the transaction's distributive implications. Importantly, when (2_w) is satisfied, A has a subjective moral reason (his belief that his gains wrong B) to modify the amount that he gains from B. Furthermore, this subjective reason is linked to an *objective* moral reason to modify his behaviour, namely, the fact that the transaction is unfair. A's failure to respond to these reasons exhibits a disrespect for, and further wrong to B that is not reducible to the transaction's unfairness. The epistemic conception is a morally distinct subset of the purely distributive conception of exploitation.

Paradigmatic cases of exploitation are *usually* cases in which the exploiters satisfy condition (2_w). We *do* have reason to believe that the distributions involved in sweatshop labour wrong sweatshop workers specifically. This epistemic condition explains the *moral* difference between 'ordinary, day-to-day transactions (that might nevertheless be unfair) and exploitations. And thus, it explains why we are not all exploiters: I do not have special reason to believe that the gains *I* receive from, say, my last bus ride to work wronged the bus company in particular. Nor does the company have reason to believe *their* gains wronged me specifically. This is so even if the pervasiveness of unfairness means that we both have a reason to believe that one of us likely gained unfairly. Most ordinary transactions are like my bus fare. They may be unfair, but we don't have evidence that one party rather than the other is the beneficiary.

However, there is a significant exception to consider: if the exploiter is *negligent*, then condition (2_w) will fail to be satisfied. Slavery represents perhaps the strongest counterexample to an epistemic conception of exploitation that appeals to (2_w). Before emancipation movements gained momentum, some slaveholders' false beliefs about a 'natural' hierarchy of races might have meant that they did not believe their transactions with slaves wronged them. Yet, surely slavery is a paradigm case of

exploitation that differs in kind from the phenomenon picked out by purely distributive accounts of exploitation. For example, it has *blameworthiness* in common with cases like sweatshops or price gouging. After all, A is blameworthy not only if he is aware (in the sense of 2_w) that his gains wrong B, but also if he is *culpably unaware* that he does so.

Yet, note that cases involving negligence must still appeal to an epistemic condition, for even if the exploiter failed to satisfy (2_w), it is the fact that he *ought* to have formed the belief that his gains wrong the other that both makes it appropriate to blame him and that distinguishes these cases from more pedestrian kinds of unfairness. Thus, according to the full version of my account, A exploits B under the epistemic conception iff

- | | |
|-----------|--|
| (1) | A gains unfairly from B and either |
| (2_w) | A believes that the gains he receives in the transaction wrong B, or |
| (2_n) | A is culpably unaware that the gains he receives in the transaction wrong B. |

The full account comprises two main conditions (unfairness and either 2_w or 2_n), each of which represents a distinctive kind of wrong. The wrong of unfairness, is the *actus reus*. The wrong of A's failure to respond to a moral reason to constrain his gain (2_w) is one possible *mens rea* wrong, but negligence (2_n) can also result in a *mens rea* wrong involving the *absence* of a mental state. The *mens rea* wrong of negligence differs from, and is arguably less bad than the *mens rea* wrong of awareness, but it is nevertheless sufficient to distinguish the epistemic conception from purely distributive approaches.

Finally, consider a possible objection: a slave owner who believes slavery is not wrong, but has doubts. He studies the issue, but everything he reads is ideological and further bolsters his false beliefs. Suppose he is neither aware, nor culpably ignorant that what he does wrongs the slave. Thus, he satisfies neither (2_w) nor (2_n) and so, does not exploit according to the account I've defended.⁸ Here I am happy to (at least partly) bite the bullet. If the slave owner is really epistemically innocent, then in this particular instance, slavery involves only an *actus reus* wrong and is not exploitative. However, I would hasten point out that such cases must have been extremely rare. *Standard* cases of slavery were exploitative because slave owners were aware, or culpably unaware that slavery wronged slaves. The counterexample does not undermine the claim that slavery was nearly always exploitative. In addition, the wrong of slavery is not exhausted by exploitation. The other wrongs of slavery still occur even if in rare cases it is not correctly described as exploitation.

I think the case is motivated in part by the worry that would-be exploiters can avoid exploitation by ensuring that they fail to satisfy one of the epistemic conditions. In practice this is very hard to achieve. Most attempts to avoid (2_w) by keeping oneself epistemically benighted will straightforwardly satisfy (2_n). Furthermore, although successful avoidance of *both* conditions would suffice to get would-be exploiters off of the conceptual hook for exploitation, it does not get them off the *moral* hook for wrongdoing. The *actus reus* wrong remains.

So, are we all exploiters? If exploitation is—as purely distributive accounts claim—tantamount to unfair transaction, then, since unfair transactions are pervasive, it seems we are. But this conclusion is not only at odds with the scope of ordinary use, it also cannot identify what is morally distinctive about paradigmatic cases of exploitation. I argued that in contrast to purely distributive accounts, an epistemic conception of exploitation according to which exploitation involves not only unfair transaction, but also the satisfaction of the epistemic condition for moral responsibility can avoid these problems while also identifying a feature—the exploiter's awareness that he wrongs his victim—that makes behaviour in paradigmatic cases so egregious.

⁸I thank an anonymous referee for this counterexample.

The epistemic conditions are commonly unsatisfied in ‘ordinary’ transactions: because we often do not know (or have any reasonable way of determining) who is taking advantage of whom, we cannot have a true belief that *our* gains in particular wrong the other. This is so even though the pervasiveness of unfairness gives us reason to believe *one of us* probably gains unfairly. Because the epistemic conditions are often unsatisfied, exploitation under the epistemic conception is less common than under purely distributive accounts, which requires only that the transaction is unfair. Therefore, unlike purely distributive accounts, my approach does not imply that we are all exploiters.

REFERENCES

- Arneson, R. (2016). Exploitation, domination, competitive markets, and unfair division. *The Southern Journal of Philosophy*, 54, 9–30.
- Arneson, R. (2013). Exploitation and outcome. *Politics, Philosophy, and Economics*, 12, 392–412.
- Arnold, D., & Bowie, N. (2003). Sweatshops and respect for persons. *Business Ethics Quarterly*, 13, 221–242.
- Bajaj, S. (2015). On Hillel Steiner’s ‘A liberal theory of exploitation’. *Ethics*, 125, 1157–1159.
- Brazier, M. (2008). Exploitation and enrichment: The paradox of medical experimentation. *Journal of Medical Ethics*, 34, 180–183.
- Cohen, G. (1979). The labour theory of value and the concept of exploitation. *Philosophy & Public Affairs*, 8, 338–360.
- Elster, J. (1982). Roemer versus Roemer: A comment on “New directions in the Marxian theory of exploitation and class”. *Politics and Society*, 11, 363–373.
- Elster, J. (1986). *Karl Marx: A Reader*. Cambridge: Cambridge University Press.
- Ferguson, B. (2016). Exploitation and disadvantage. *Economics and Philosophy*, 32, 485–509.
- Ferguson, B. (2018). Exploitation. In William Thompson (Ed.), *Oxford Research Encyclopaedia of Politics*. New York: Oxford University Press.
- Ferguson, B., & Steiner, H. (2018). Exploitation. In S. Olsaretti (Ed.), *The Oxford Handbook of Distributive Justice* (pp. 533–555). Oxford: Oxford University Press.
- Ferguson, B., & Ostmann, F. (2018). Sweatshops and consumer choices. *Economics and Philosophy*, 34, 295–315.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Geras, N. (1985). The controversy about Marx and justice. *New Left Review*, 150, 37–69.
- Goodin, R. (1987). Exploiting a situation and exploiting a person. In A. Reeve (Ed.), *Modern theories of exploitation* (pp. 166–197). London: Sage.
- Kalai, E., & Smorodinsky, M. (1975). Other solutions to Nash’s bargaining problem. *Econometrica*, 43, 513–518.
- List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Marx, K. (1867). *Capital: A Critique of Political Economy*, volumes 1–3. Chicago: C.H. Kerr and Company.
- Mele, A. (2010). Moral responsibility for actions: Epistemic and freedom conditions. *Philosophical Explorations*, 13, 101–111.
- Nash, J. (1950). The bargaining problem. *Econometrica*, 18, 155–162.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Oxford: Blackwell.
- Olsaretti, S. (2009). Responsibility and the consequences of choice. *Proceedings of the Aristotelian Society*, 109, 165–188.
- Ostmann, F. (2017). *Fairness in International Commercial Transactions*. Doctoral thesis, University College London.
- Peels, R. (2014). What kind of ignorance excuses? Two neglected issues. *Philosophical Quarterly*, 64, 478–496.
- Powell, B., & Zwolinski, M. (2011). Ethical and economic case against sweatshop labor: A critical assessment. *Journal of Business Ethics*, 107, 449–472.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Reiff, M. (2013). *Exploitation and Economic Justice in the Liberal Capitalist State*. Oxford: Oxford University Press.
- Roemer, J. (1982a). *A General Theory of Exploitation and Class*. Cambridge, MA: Harvard University Press.
- Roemer, J. (1982b). New directions in the Marxian theory of exploitation and class. *Politics and Society*, 11, 253–287.
- Rosen, G. (2008). Kleinbart the oblivious and other tales of ignorance and responsibility. *Journal of Philosophy*, 105, 591–610.

- Rudy-Hiller, F. (2018). The epistemic condition for moral responsibility. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition).
- Robichaud, P., & Wieland, J. (2017). *Responsibility: the Epistemic Condition*. Oxford: Oxford University Press.
- Sample, R. (2003). *Exploitation: What It is and Why It's Wrong*. Boulder, CO: Rowman and Littlefield.
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, 1, 229–243.
- Sliwa, P. (2017). On knowing what's right and being responsible for it. In P. Robichaud & J. Wieland (Eds.), *Responsibility: the Epistemic Condition* (pp. 127–145). Oxford: Oxford University Press.
- Snyder, J. (2008). Needs exploitation. *Ethical Theory and Moral Practice*, 11, 389–405.
- Steiner, H. (1984). A liberal theory of exploitation. *Ethics*, 92, 225–241.
- Steiner, H. (1987). Exploitation: A liberal theory amended, defended and extended. In A. Reeve (Ed.), *Modern Theories of Exploitation* (pp. 132–48). London: Sage.
- Steiner, H. (2010). Exploitation takes time. In J. Vint, J. Metcalfe, H. Kurz, N. Salvadori & P. Samuelson (Eds.), *Economic Theory and Economic Thought: Essays in Honour of Ian Steedman* (pp.20–29). London: Routledge.
- Steiner, H. (2018). Exploitation, intentionality, and injustice. *Economics and Philosophy*, 34, 369–379.
- Vrousalis, N. (2013). Exploitation, vulnerability, and social domination. *Philosophy and Public Affairs*, 41, 131–157.
- Vrousalis, N. (2016). Exploitation as domination: A response to Arneson. *Southern Journal of Philosophy*, 54, 527–538.
- Walt, S. (1984). Comment on Steiner's "Liberal theory of exploitation". *Ethics*, 92, 242–247.
- Wertheimer, A. (1996). *Exploitation*. Princeton: Princeton University Press.
- Wieland, J. (2017). Introduction: The Epistemic Condition. In P. Robichaud & J. Wieland (Eds.), *Responsibility: The Epistemic Condition*, (pp. 1–28). Oxford: Oxford University Press.
- Wilkinson, S. (2003). The Exploitation argument against commercial surrogacy. *Bioethics*, 17, 169–187.
- Wolff, J. (1999). Marx and exploitation. *The Journal of Ethics*, 3, 105–120.
- Wood A. (1972). The Marxian critique of justice. *Philosophy & Public Affairs*, 1, 244–282.
- Wood, A. (1981). Marx and equality. In J. Mepham & D. Hillel-Ruben (Eds.), *Issues in Marxist Philosophy* Vol. 4. (pp. 195–221). Sussex: Harvester Press.
- Wood, A. (1995). Exploitation. *Social Philosophy and Policy*, 12, 135–158.
- Wood, A. (2016). Unjust exploitation. *The Southern Journal of Philosophy*, 54, 92–108.
- Zimmerman, M. (2017). Ignorance as a moral excuse. In R. Peels, (Ed.), *Perspectives on Ignorance from Moral and Social Philosophy* (pp. 77–94). New York: Routledge.
- Zimmerman, M. (2002). Controlling ignorance: A bitter truth. *Journal of Social Philosophy*, 33, 483–490.
- Zwolinski, M. (2008). The ethics of price gouging. *Business Ethics Quarterly*, 18, 347–378.
- Zwolinski, M., & Wertheimer, A. (2016). Exploitation. In Edward N. Zalta (Ed.), *The Stanford Encyclopaedia of Philosophy* (Fall 2016 Edition). <https://plato.stanford.edu/archives/fall2016/entries/exploitation/>.
- Zwolinski, M. (2012). Structural exploitation. *Social Philosophy and Policy*, 29, 154–179.

How to cite this article: Ferguson B. Are we all exploiters?. *Philos Phenomenol Res.* 2020;00:1–12. <https://doi.org/10.1111/phpr.12730>