




ARTICLE

Exploitation's grounding problem

Benjamin Ferguson 

Department of Philosophy, University of Warwick, Coventry CV4 7AL, UK
Email: benjamin.ferguson@warwick.ac.uk

(Received 01 February 2024; revised 11 September 2024; accepted 16 September 2024)

Abstract

Standard accounts of what makes exploitation wrong ground its wrong in distributive unfairness: when A exploits B he wrongs her by taking a greater share of the benefits from their interaction than he ought. I argue that this standard account does not succeed; distributive unfairness is neither the sole, nor the primary wrong of exploitation. I assume that distributive unfairness is pro tanto wrong. However, I argue that in situations where transactors' consent to a transaction is morally valid, it is also morally transformative and overrides distributive unfairness's pro tanto wrong. Thus, wrongful exploitation requires morally invalid consent.

Keywords: Exploitation; consent; unfairness; distributive justice

1. Introduction

Suppose B is seeking a job, A is hiring workers, and A offers B a job with an unfairly low wage.¹ Though unfair, this wage is better than B's next best option – perhaps it provides a better deal than her next best job offer, or perhaps it is better than her status quo. Since A's offer is better than B's next best option it benefits B compared with that option. Consequently, B accepts A's offer.

Most philosophers agree that distributive unfairness makes, or is at least part of what makes, a transaction an exploitation, though, naturally, there is some disagreement about the precise analysis of exploitation.² This paper is not about the scope of (transactional) exploitation. It is about what makes it wrong.

¹In what follows I shall be neutral about what fairness amounts to in the context of transactions. Readers should feel free to substitute their own accounts of transactional fairness.

²Here, for simplicity, I shall equate unfair transaction and exploitation. There are many problems with this characterization of exploitation (Ferguson 2021; Zwolinski *et al.* 2022). Most of these issues do not affect the attempt to ground exploitation's wrong on its unfairness. Note however, that in this paper my focus is on *transactional* exploitation, which takes place in a discrete transaction, with an identifiable exploiter and an identifiable victim. I should note that my arguments do not straightforwardly apply to systemic accounts that identify exploitation with the aggregate distribution of values in an economy (Morishima 1974; Roemer 1982; Yoshihara and Veneziani 2018).

A common view is that sweatshops are wrong because they involve unfair pay.³ Professional philosophers widely agree. Most accounts of exploitation claim unfairness grounds exploitation's wrong.⁴ Many accounts of exploitation are purely distributive accounts that accept the claim that 'wrongful exploitation is tantamount to unfair transaction' (Ferguson 2021: 536). Further, those accounts that appeal to other wrong-making features often cash these wrongs out in terms of distributive unfairness (Wolff 1999; Sample 2003).

I will argue that this common position cannot be sustained. If mutually beneficial exploitations *are* wrong, this wrong cannot be grounded *solely* by the transaction's distributive unfairness. The reason is that if the exploited party offers morally valid consent to the transaction, this consent morally transforms the pro tanto wrong of unfairness, making the transaction pro tanto permissible. Thus, distributive unfairness is insufficient for wrongful exploitation. Or so I shall argue. Call this the 'grounding problem'.

Further, I will show that distributive unfairness is not only insufficient for wrongful exploitation, it is also morally and causally subsidiary to the kinds of wrongs that invalidate consent and these wrongs do not include distributive unfairness. The handful of accounts that posit wrong-making conditions that are irreducible to distributive unfairness (Ferguson 2021) nevertheless fail to mention this primacy of invalid consent (if, indeed, they include it at all).

Finally, I will argue that in many paradigmatic cases of exploitation there is good reason to believe the transactors' consent is morally valid. If morally valid consent is also morally transformative, then this means that many paradigmatic cases of exploitation are not, in fact, morally wrong.⁵

One way to argue that exploitations are not wrong is by appealing to what has variously been called the 'paradox of exploitation' (Ferguson 2016; Faraci 2019), the 'exploitation problem' (Horton 2019), and the 'nonworseness problem' (Wertheimer 1996; Zwolinski 2009). Assume that it *is* morally permissible for A to leave B with her next best option. If (1) A's offer is *morally* better than this option because it provides B a benefit compared with her next best option, and if (2) acts

³Or, more accurately, an unfair remuneration package of pay and working conditions.

⁴The idea that exploitation is wrong because of distributive unfairness is endorsed by John Roemer (1982), Alan Wertheimer (1996: 13), Ruth Sample (2003: 57), Mark Reiff (2013: 39), Hillel Steiner (2013: 338) and Richard Arneson (2016: 12) (see also Steiner 1984, 2010), and nicely summarized in papers by Jonathan Wolff (1999) and by Eric Malmqvist and András Szigeti (2021). Benjamin Ferguson (2013, 2021) argues that unfairness explains the *actus reus* wrong of exploitation, but locates the *mens rea* wrong in exploiters' mental states. Robert Goodin (1987) and Mikhail Valdman (2009) both argue that exploitation is wrong because it involves the extraction of excessive benefits from the vulnerable, analyses that most straightforwardly fall into the distributive unfairness camp. Two notable exceptions to the claim come from Alan Wood (1995) and Nicholas Vrousalis (2013). Vrousalis, in particular, argues 'the wrongmaking features of [exploitation] must be sought neither in coercion, nor in force, nor . . . in distributive injustice' Rather, for Vrousalis, exploitations are wrong because they involve domination.

⁵Whether such cases should then be labelled 'exploitations' is a matter of debate because theorists disagree about whether exploitation is necessarily moralized (see Wood (1995) for an argument that it is not). If exploitation is necessarily moralized, then would-be exploitations that are not wrong are not exploitations and might best be referred to as 'unfair transactions'. If exploitation is not necessarily moralized we can still refer to such cases as (permissible) exploitations. For simplicity, I will opt for the latter terminology, with the caveat that this does not imply I am taking a substantive stance in the debate.

that are morally better than morally permissible options are also morally permissible, then an agreement between the two regarding *any* welfare improving offer must be morally permissible, even if the distribution of benefits is unfair. Or, in other words, given (1) and (2), unfair employment is *impermissible* only if leaving B with her next best option is also impermissible. Much has been written about this argument (Zwolinski 2009; Ferguson 2016; Faraci 2019; Horton 2019; Berkey 2021). If the paradox argument is sound, then it frustrates attempts to show that unfair, but mutually beneficial transactions (and perhaps, by extension, exploitations) can be morally impermissible. I think many philosophers believe that if the paradox argument fails, it does so by way of condition (1) which requires accepting a weak form of welfarism – a position that some reject (Sen 1979).

In this paper I will present a different argument that undermines the claim that many exploitations are wrong. This argument does not depend on the truth of claims (1) and (2) from the paradox argument, but instead appeals to the morally transformative power of consent and transactors' property rights over their own labour and goods. However, there are different ways that mutual benefit and the permissibility of non-interaction intersect with the grounding problem which I will revisit below.

I shall assume that distributively unfair transactions are pro tanto impermissible. Thus, if the terms of an employment contract between A and B are unfair, then entering into the contract is pro tanto impermissible. Of course, acts that are pro tanto impermissible might not be all things considered impermissible if there are other factors at play that mitigate or transform this impermissibility. But when a transaction is distributively unfair to B, she has at least a pro tanto moral claim to better terms and, I shall assume, she is (again, pro tanto) wronged by A when A impermissibly transacts with her on unfair terms. What, precisely, the terms of a permissible transaction that does not wrong B should be will depend on the account of fairness one adopts.

Suppose, as is true in the sweatshop case outlined above, that B assents to an unfair employment offer from A. Some might find it puzzling that she accepts unfair terms when she has an apparent moral claim to better ones. But, of course, as Marx (1867) understood, in practice, workers accept these terms because they lack the power to obtain better ones. Two crucial questions are (1) whether and when B's assent to unfair terms can be described as morally valid *consent*, and (2) whether in such cases morally valid consent is sufficient on its own to make what would have been an impermissible transaction morally permissible.

There is a third question that I will not cover here: given valid consent, is *third-party interference* permissible? Matt Zwolinski (2007), Jessica Flanigan (2018) and David Faraci (2019) have argued that it is not; Michael Kates (2015) argues it is. Whether it is or not depends, I think, on its collective effect on workers' welfare. More importantly though, my claim in this paper is not about third-party interference. It is both different and stronger: morally valid consent is morally transformative in the transaction itself, overriding exploitation's pro tanto wrong.

In the rest of the paper, I tackle the first two questions in reverse order. In sections 2 and 3 I consider whether morally valid consent is morally transformative. Section 2 considers the question theoretically, asking whether B has the moral power to waive claims she has to a fair transaction. Section 3 takes an applied approach.

There I argue that we assume morally valid consent is morally transformative in cases that are relevantly similar to exploitations. After arguing that valid consent to unfair transactions is morally transformative, I briefly revisit the relationship between my argument and the paradox of exploitation in section 4 before turning to the question of whether and when B's consent is morally invalid. Sections 5 through 8 consider four features that might undermine the validity of consent. In section 5 I discuss fraud, section 6 pressure, section 7 coercion, and in section 8 distributive unfairness. In sections 9 and 10 I briefly discuss objections that appeal to externalities and partially valid consent, respectively, before concluding with section 11.

2. Transformative Consent in Theory

In many situations, acts that would be seriously wrong are morally transformed by valid consent. Consent is crucial in the transformation of theft into borrowing, rape into permissible intercourse, and assault into a boxing match. Is morally valid consent also sufficient to morally transform the impermissibility of mutually beneficial unfair transactions? Matt Zwolinski (2007: 696) has argued that workers' choices to accept the conditions of sweatshop labour contracts 'establishes a claim of non-interference against those who might wish to prevent them from engaging in sweatshop labor', but he notes that his argument 'does not even attempt to show that their choice renders such treatment morally permissible'. I will argue that when these choices involve morally valid consent they *do* transform moral permissions.

To begin, let us consider what kind of rights B has to her own labour. Of particular importance is the scope of B's *second-order* right, or power, to transform her first-order right of self ownership.⁶ To say that B has the moral power to rent her labour in the context of an unfair transaction is to say that her valid consent to enter into these contracts is morally transformative.

It is uncontroversial that in the context of working for A in a sweatshop (and, indeed, in most⁷ cases) B has a first-order right of self-ownership that implies a claim right over her own labour, and, in turn, a correlative duty for A to refrain from enslaving her. It is also uncontroversial that B typically has the power to transform this first-order right by allowing others the use of her labour. What *is* controversial is the *scope* of this power. Libertarians who endorse full liberal ownership argue that B's power includes even the ability to sell herself into slavery, effectively a third-order right to transfer her power over her claim right to another person. A more widely accepted view is that B has the power to *rent* her labour to A, but this power itself is inalienable. That is, B *lacks* a moral right to the permanent third-order transfer right. Defences of this view typically appeal to claims that persons are not the kinds of things that can be owned.⁸

⁶Here and throughout I will adopt the widely accepted Hohfeldian system of rights (Hohfeld 1919).

⁷The scope of B's claim right to her own labour is plausibly less than universal. B may, for example, have an obligation to aid A in certain circumstances. But in sweatshops and other exploitation-adjacent situations, B has no such obligation to A. See Cohen (1995: Ch. 10) for a discussion.

⁸Similar arguments have been used to argue against the permissibility of commodifying organs, sex, votes, and other goods.

Crucial for our purposes is whether there are any additional limits on B's power that constrain the *terms* she can accept for the use of her labour. If there are, then even if B offers morally valid consent to an unfair transaction, she may nevertheless lack the moral power to engage in it, which would mean her valid consent is nevertheless *not* morally transformative and, consequently, A's employment of B on unfair terms remains morally impermissible.

So, are there limits to B's power to rent her labour? There are two reasons for scepticism. The first, which I will discuss in this section, is that in the context of unfair transactions, limits on B's powers are not supported by either of the most prominent justifications of rights – the interest theory and the will theory. The second, discussed in section 3, is that we do not assume limits to B's power in similar cases.

First consider the will theory of rights. According to the will theory (Hart 1982; Steiner 1994), the idea that B could lack the power to waive her right to a fair transaction is a non-starter. For will theorists, B holds a right if and only if she 'is competent and authorised to demand or waive the enforcement of the right' (Kramer *et al.* 1998: 62). Rights make their holders 'small-scale sovereigns' and the power to waive one's rights is built into the definition of rights (Hart 1982: 183). For will theorists, if B has a right against A using her labour, then she has the power to waive this right. On the will theory then, B's morally valid consent is sufficient to morally transform the pro tanto impermissibility of unfair transaction.

Now consider interest theories.⁹ Interest theories claim a necessary condition for B holding a right is that the right protects her interests. In contrast to will theories, interest theories deny that being competent and authorized to demand or waive the holding of a right are necessary and sufficient conditions for holding a right (Kramer *et al.* 1998: 62). It is typically in B's interests to have the power to rent her labour to others, so on the interest theory she usually has such a power. But we can also ask whether, more narrowly, it is in her interest to have this power in circumstances where the transaction is unfair, for even if she typically has the power to waive rights over her labour, the scope of this broader right could be constrained if it doesn't serve her interests in particular cases.

One problem for interest theorists is that if a transaction is unfair but nevertheless mutually beneficial, it *does* further B's interests. It leaves her better off than she would have been had she not transacted. This suggests that the interest theory entails B has the power to rent her labour even in unfair transactions.

However, interest theorists could respond that it is also in B's interests to be treated fairly in transactions. Consequently, she has another right: a right against A transacting unfairly with her. If the interest theory also entails that B lacks the power to waive *this* right, then when A transacts unfairly with her A violates B's right to a fair transaction and this violation can ground exploitation's wrongfulness. Kramer *et al.* (1998: 78) suggest that the interest theorists can make this move, arguing that the theory 'can attribute rights to workers under a minimum-wage law even though the workers cannot themselves waive or enforce their entitlements to be paid at a certain level'.

⁹I remain agnostic here about whether interests (and benefits) should be defined subjectively, objectively or in a hybrid way. The only requirement is that they are defined consistently between interests and benefits.

Yet, this conclusion might be too quick. The problem for the interest theory is that in many cases, A has no general duty to transact with B. Rather, obligations to not exploit, e.g. obligations to transact fairly, are understood to be *conditional* obligations, such that *if* A transacts with B, *then* A ought to do so on fair terms. In other words, A's obligation to transact fairly with B applies only when, and in virtue of the fact that, A transacts with B (Wertheimer 1996, 2010; Sample 2003; Ferguson 2016; Rulli 2020). To return to rights language, this means A has a *liberty* to not transact with B. And if A has such a liberty, then B has no right for A to transact with her. The problem for the interest theorist is that if protecting B's interests does not generate a general right to be transacted with, and the protection of A's interests does grant him a liberty to not transact, then it is hard to see how B's interests can be insufficiently protected by an unfair, but mutually beneficial transaction that leaves her *better off* than not transacting.¹⁰ Or, to put the point the other way around, if the protection of B's interests would justify a right to a fair transaction, then it seems it would also justify a right against not being transacted with at all since not being transacted with at all serves her interests worse than an unfair transaction. The ascription of rights in exploitation cases that we might find *intuitive* is nevertheless non-monotonic with respect to interests. We would *like* to say that B has a right against unfair transaction, but not a right against non-transaction. However, B's interests are promoted more by fair transaction than unfair transaction and more by unfair transaction than non-transaction. Interest theorists cannot have it both ways. Either A has no duty to transact with B and B has no right to a fair transaction, or B has a right to a fair transaction *and* A has a general duty to transact with her. Although this latter option may be plausible in certain cases, such as those in which B is in duress, positing a general obligation to engage in mutually beneficial transactions is not only at odds with the standard assumption that duties to not exploit are conditional, it is also incredibly morally demanding.

So, to sum up: the will theory of rights is incompatible with the claim that B has a right to a fair transaction but lacks the power to waive it. The interest theory fares better. It *is* in B's interest to obtain a fair price rather than an unfair price. But the problem for interest theories is that it is also in her interest to obtain an unfair, but mutually beneficial price rather than to not transact at all. 'Pure' interest and will theories do not exhaust existing theoretical approaches. A number of hybrid theories have emerged in recent years (Sreenivasan 2005). However, while these theories solve some traditional problems confronting the will and interest theories, it is unclear how they can be used to support the claim that valid consent to an unfair transaction is not morally transformative.

Finally, note that my theoretical conclusions here should not be taken to imply there are *no* inalienable rights. Some rights may be inalienable; for example, Terrance McConnell (2000) defends the inalienability of the right to life, the right of conscience, and rights against non-consensual medical treatment. My claim here is

¹⁰This point appeals to the non-worseness claim that generates the paradox of exploitation. In a sense then, non-worseness can be the source of both the grounding problem and the paradox of exploitation for interest theorists, even though those two problems remain distinct. While some theorists solve the paradox by rejecting non-worseness it is unclear that this solution is open to the interest theorist since their theory relies precisely on interests, e.g. those things that are good for the agent in question.

more limited: neither the interest nor the will theories support the claim that the *right to a fair transaction* is an inalienable right. In the next section I will provide additional practical support for this claim by showing that we take valid consent to be morally transformative in other cases that are relevantly similar to exploitation.

3. Transformative Consent in Practice

It is widely accepted that morally valid consent is morally transformative in cases relevantly similar to exploitations: harmful transactions and donations.

First, consider harmful transactions. The consumption of tobacco products, the purchase of lottery tickets, risky investments, and many other goods tend to make consumers worse off than they would have been had they not purchased these items. Yet, provided consumers' consent is morally valid, it is sufficient to waive any claims they might have against harms accompanying these purchases. There are, of course, many circumstances associated with harmful transactions that might *invalidate* transactor's consent. Indeed, a transactor's enthusiasm for an apparently harmful transaction will probably make us suspicious that they have been subject to misleading advertising or that they do not really understand what they are getting themselves into. But if these invalidating factors are absent and the transactor really does provide morally valid consent to the transaction, then this consent is broadly understood to be morally transformative as well.

Unlike harmful transactions, mutually beneficial but unfair transactions do not make transactors worse off than they were before transacting. If valid consent is morally transformative for harmful transactions, then surely it can also morally transform mutually beneficial exploitations.

Second, consider donations. It is widely assumed that B's valid consent makes it permissible for her to *donate* her labour to A for free. If B's valid consent is morally transformative in a transaction involving a fair distribution of economic benefits *and* it is morally transformative at a price in which A receives all of the economic gain, then it seems that it should also be transformative at a price between these points where the gain is unfairly distributed.

As an illustration, consider the following case. This past year I received a new smartwatch and decided to sell my old one online. I was disappointed to discover that it would probably only fetch \$75 on eBay. Apparently technology moves on quickly. Let's suppose for the sake of the argument that \$75 is a fair price. I remembered my neighbour wanted a smartwatch and decided to offer mine to her for \$20. My partner pointed out that \$20 seemed a little avaricious – if I was going to offer it to the neighbour, I should just *give* it to her. Of course, she was right and I gave my neighbour the watch.

My property ownership of the watch also entailed that I had the power to accept whatever price I wanted in exchange for it – \$75, or \$20, or nothing at all. If this is true of the property rights I have over my possessions, it is, I submit, also true of the property rights I have over my labour. If I can give my labour away and I can sell it for a fair price, then – provided my consent is morally valid – my choice to sell it for an unfair price morally transforms any claims I might have had against being transacted with unfairly.

Now, an interesting rebuttal here is that two transactions can be equivalent in their terms while nevertheless being different kinds of transaction. Had I exchanged the watch for \$20, this would have been what Steiner (1984: 226) has called a 'benefit'. Benefits and unfair transactions can have the same terms, but benefits differ from unfair transactions because in a benefit I would not accept more than \$20 for the watch. Doing so would defeat my purpose of quasi-donating it to my neighbour since my intention was to give her most (but not all) of the economic gain. On the other hand, in an *unfair* transaction, a transactor hopes to get as much as possible for the goods they exchange. Unlike in a benefit, transactors would happily accept a greater share of the economic gains in unfair transactions. Thus, although they may have identical terms – a watch exchanged for \$20 – benefits and unfair transactions can be distinguished counterfactually, by appealing to transactors' attitudes towards gain. Therefore, although donations and benefits might be morally permissible when consent is morally valid, the differences between these transactions and unfair transactions means that valid consent is *not* transformative for unfair transactions.

The distinction between donations/benefits and unfair transactions is well taken. But the rebuttal doesn't do the work required to undermine my claim. In fact, it adds credibility to the idea that valid consent is morally transformative, for it suggests that an unfair transaction can *become* a benefit at the discretion of the transactor. Whether or not a transaction is a benefit, a donation, or an unfair transaction is, in an important sense, up to the transactor since the difference between these transaction types depends on her attitudes. This means that morally valid consent *is* morally transformative in some transactions that would otherwise have been unfair: in the watch example, \$20 would have been an unfair transaction until *I* opted for a benefit (and eventually, a donation). The point is that whether a transaction is a donation, benefit, or unfair transaction is at the discretion of the transactor.

To be clear, there is a difference between B's choosing between different transaction *types*, which is a second-order kind of choice, and her consenting to particular terms within a type. But it is not clear why this difference should make a *moral* difference. Consider the watch case again. So far we have established that if I either donate the watch or exchange it as a benefit, my valid consent to enter into these types of transaction *is* morally transformative. Note that in both of these transaction types, I can forego the economic gains I could have received in a standard transaction for other regarding reasons. Further, I can choose between these types of transaction and a standard transaction (selling the watch on eBay). But, if my valid consent cannot transform unfair transactions, then, if I decide at the second-order level to engage in a standard transaction and maximize my own gains, and I decide that validly consenting to an unfair transaction is the best way of doing so, the transaction remains impermissible. Why should transactors have the power to forego economic gains for other regarding reasons, but not to engage in unfair transactions for self-regarding reasons?

Further, maintaining that valid consent transforms donations and benefits but not unfair transactions means that whether A wrongs B by impermissibly engaging in an unfair transaction with her depends on B's mental state. Since it is B's views about what counterfactual terms she would accept that distinguish a benefit from an

unfair transaction and not the terms of the transaction, and since B's attitudes can be inaccessible to A, in many cases whether A wrongs B is a matter of moral luck.

Note that the mental state in question is not something like B's having a pro attitude towards transacting. There is nothing odd about that kind of mental state making a difference to moral permissions, precisely because B's wanting (or not wanting) a particular action to take place can be communicated to A by B consenting. B can hold pro attitudes towards transacting in both unfair transactions and benefits (and she can communicate these attitudes via consent in both cases). If we are to treat the two as morally different, this means that two transactions can be identical in their terms, that B can validly consent to both and have pro attitudes towards both, and yet, because B would have been willing to accept more in the unfair transaction, the transactions' moral statuses differ. It is this last condition – B's attitude about *counterfactual* terms – that may reasonably be unknown to A and that generates moral luck.

Neither theoretical nor practical considerations undermine the claim that valid consent is transformative in unfair transactions. I suspect most readers think – as I do myself – that the action really lies in the discussion of conditions that characterize morally valid consent. If you think that transactors' consent in exploitations is typically *invalid*, then the question about the transformative power of valid consent seems otiose. I will discuss morally valid consent in sections 5 through 8. But, first, note that the conclusion that valid consent is morally transformative is a meaningful departure from many accounts of what makes exploitation wrong. We can no longer simply infer from the unfairness of a transaction that it wrongs the party on the sharp end of the exchange, for if this party *did* provide valid consent, they would no longer be wronged. Therefore, an unfair distribution of benefits is not a sufficient condition for wrongful exploitation. This result is contrary to many existing accounts of exploitation that analyse the wrong of exploitation in purely distributive terms. If morally valid consent is morally transformative in unfair transactions, purely distributive accounts cannot be correct. The crucial question now is when is consent morally *valid*. However, before addressing this question I will briefly revisit the paradox of exploitation.

4. The Paradox Revisited

In the previous section I argued that provided B's consent is valid, this valid consent morally transforms the pro tanto wrong of unfair transaction. Appeals to B's moral right to a fair transaction – whether will-based or interest-based – are not successful. Perhaps more importantly, we commonly recognize the morally transformative power of B's consent even in situations where she engages in harmful transactions.

Note that the paradox of exploitation argument does not depend on consent, and my argument here does not necessarily depend on conditions in the paradox like nonworseness or the permissibility of non-interaction. However, these conditions can play contingent role in my arguments in two ways.

First, if one is a pure subjectivist about welfare then the fact that B consents is (almost) all that is needed to make the transaction beneficial, since 'better than' for pure subjectivists is just the satisfaction of preferences. So unless B is mistaken about

her preferences, her consent makes it such that the transaction is mutually beneficial (because the same holds for A as well). Thus, for pure subjectivists consent and mutual benefit can coincide, partially linking the arguments.

However, consent and welfare improvement needn't coincide because there are other accounts of welfare. Welfare and consent are conceptually distinct and my arguments apply not only to purely subjective accounts, but also to objective list and hedonic accounts. So, though the arguments may overlap on this point for pure subjectivists, they are nevertheless conceptually distinct.

Second, my reply to the interest theorist's counterargument appeals to both the claim that mutually beneficial transactions promote the transactors' interests and to the claim that there is a right against non-transaction. These claims also feature in the paradox.

However, their role is to explain why the interest theorist's counterargument cannot succeed: interest theory isn't capable of condemning mutually beneficial exploitation if it also doesn't recognize a right against non-transaction. This is because rights kick in at a certain level of interest in a monotonic fashion – if the promotion of interests at level x generates a right, then so does the promotion of interests at level $x + n$. Appeals to the permissibility of non-interaction and mutual benefit are not part of the positive argument for the grounding problem, they are not necessary for showing that will theorists cannot escape the problem and they are not necessary for showing that we take consent to be transformative in similar contexts.

Finally, it is worth pointing out that my arguments can be used to make a claim that is far broader in scope than the paradox of exploitation: if one thinks that valid consent can be transformative even in the presence of objective harm (such as in the case of, say, cigarette sales) then one might argue that morally valid consent can transform the wrong not only of mutually beneficial exploitations, but also harmful exploitations. I have stopped short of making that argument here, for it is a controversial and far stronger claim, which becomes more controversial the greater the harm. But if we think that valid consent can at least be transformative in the presence of minor harms, then my arguments could be employed to argue that some instances of harmful exploitation are not morally wrong in the presence of morally valid consent.

5. Valid Consent and Fraud

It is well-known that not all *assents* actually amount to morally valid forms of *consent*. Agreeing to let someone take your car at gunpoint does not involve valid consent; nor does assenting to a medical procedure without proper information about one's alternatives, or assenting to sex under the influence of drugs. Similarly, if B agrees to an unfair employment contract with A under the threat of violence from A, B has not provided morally valid consent and, consequently A still wrongs B by engaging in an unfair transaction because B's mere assent is not morally transformative.

Some sweatshops contracts described by the label 'modern slavery' are situations where workers don't even offer assent. Such cases are simply and straightforwardly instances of slavery and involve the theft of workers' labour. In other cases, workers

do assent, but their assent is predicated on deceit. Workers are lured into contracts on false pretexts and threatened physically or with (supposed) debts to their employer when they try to leave. This kind of employment is morally and conceptually akin to more straightforward forms of fraud and can become slavery when workers are prevented from leaving. If workers would not have assented to their contracts had they not been deceived, then their assent does not amount to morally valid consent (Ferguson 2018). Cases of exploitation involving fraudulent contracts are therefore morally wrong and, since there is no valid consent to transform the wrong of the transaction's unfairness, they are also wrongfully unfair.

Note, however, that the *primary* wrong in such cases is the fraud and deception. This wrong morally activates the wrong of unfairness since it undermines the transformative power of the defrauded party's consent by making the consent invalid. Absent the fraud, the consent would have been valid and, consequently, morally transformative. Note that fraud is also causally responsible for the unfairness. By lying about the transaction's terms, the exploiter is able to gain more for themselves than they otherwise would have gained had the contract's terms been truthfully communicated. Thus, not only is unfairness insufficient on its own to ground wrongful exploitation, but in cases involving fraud it is also morally and causally subsidiary to deceit.

However, not all cases that we might describe as unfair or exploitative involve fraud. Sometimes workers accept unfairly low pay despite having accurate knowledge of their contract's details. They do so because even though unfair, the contracts offer a far better package than their next best option. Employment in a sweatshop often offers better pay and conditions than ragpicking, which in turn, can be a better option for some than subsistence agriculture. Is workers' assent to contracts under these conditions morally valid consent?

6. Valid Consent and Pressure

One reason to think it is not is that workers choice is 'pressured'. Workers' choices might be pressured in two senses. First, they might face choices between two unacceptable options. We might think that consent is valid only if the options one is choosing between are sufficiently good (Olsaretti 1998). I'll return to this sense of pressure in the next section. Here I want to focus on transactions that are pressured in a second sense: they involve high stakes and/or the need to make a choice quickly.

Now, *all* mutually beneficial transactions will involve *some* degree of pressure in this second sense because the transactors have an interest in the transaction occurring since it is better than their status quo option and there will typically be some limit to the time they have to make their decision. Suppose B has just finished her PhD and has the good fortune of receiving *two* tenure track job offers, which she needs to respond to within weeks. Department X is more prestigious and involves a lighter teaching load, but department Y offers higher pay in a better location. B faces a high-pressure decision, but it is probably not one that undermines the validity of her consent. An appeal to pressure as a factor that vitiates consent requires drawing a line between mere pressure and vitiating pressure, and doing so in a non-arbitrary way. Had department Y required an instantaneous response from B, we might

conclude that *that* kind of pressure was vitiating. There are many different ways to distinguish between vitiating and mere pressure, but let's suppose that a principled distinction can be made. An appeal to pressure must also explain what, in particular, is *vitiating* about vitiating pressure. The most obvious explanation is that extreme pressure can lead to deliberative mistakes that undermine autonomy. Under pressure, B may agree to terms or choose options she later regrets.

Note that although some forms of pressure might be morally wrong – one thinks here of used car salesmen preying on vulnerable persons – pressure itself is not necessarily morally wrong. A search committee could have good fiscal or legal reasons to require a very quick response from a candidate. But even if such pressure is not introduced for nefarious reasons, it nevertheless undermines the validity of consent when it interferes with deliberation. However, unfair terms are typically tied to the nefarious use of pressure, for the point of such pressure is precisely to interfere with transactors' ability to detect and deliberate about unfair terms.

The nefarious pressure introduced to interfere with deliberation about an unfair offer is one kind of vitiating pressure that is morally and causally similar to fraud. It is the *pressure* that undermines the transformative power of consent and which enables exploiters to gain more than they otherwise would have. And as with cases involving fraud, even though the transaction's unfairness remains wrong, its wrong is subsidiary to the wrong of vitiating pressure.

Some cases of exploitation certainly involve this constellation of wrongs. Exploitations arising from used car salesman-like tactics are prime examples. But deliberation undermining pressure is not a feature of some of the most salient cases of exploitation. Suppose that B faces a choice not between an academic job at departments X or Y, but rather between an academic job at department X, sweatshop work, and subsistence farming. Suppose she prefers the job at X to sweatshop work and sweatshop work to subsistence farming. If her options are reduced by the removal of the academic job, they become much *worse*, but she is still going to prefer the sweatshop over subsistence farming. Indeed, it is those in the most precarious circumstances who will also be most acutely aware of – and insistent about obtaining – the benefits of their best available option. In these cases B is unlikely to make a deliberative mistake, accidentally choosing a job that leaves her worse off. She will choose the option that furthers her interests, even if it involves unfair terms.

These types of cases might still involve invalid consent, but not because pressure in the sense I've been discussing leads to a deliberative mistake. Rather, such cases might involve invalid consent because of the *other* sense of pressure mentioned above: being forced to choose between two morally unacceptable options.

7. Valid Consent and Coercion

One reason to think that transactors' consent is morally invalid in many paradigmatic cases of exploitation is none of the available options are morally acceptable. Another reason that their consent might be invalid is that it is coerced. Both of these situations can be analysed as situations in which coercion undermines the validity of consent. Following Nozick's (1969) and related accounts of coercion

with a moralized baseline, let us say that A coerces B iff A gets B to φ by credibly threatening to ψ if B doesn't φ and that what A threatens to do – ψ -ing – is morally impermissible.

First, consider situations from the exploitation literature in which B is saved from a snow storm, from drowning, or from dehydration by paying A an extortionate amount or performing sexual favours (Wertheimer 1996; Meyers 2004). These cases involve coercion, at least under the plausible assumption that A had an obligation to help B in the situations described. By threatening to leave B unless she pays up, A threatens to impermissibly shirk his obligation to save her. Extracting assent from her under this threat is coercive. A's coercion undermines the moral validity of B's consent.¹¹

Second, consider situations where none of B's available options are morally acceptable. Suppose these options are working in a sweatshop or ragpicking and that ragpicking is her status quo option. If neither is morally acceptable, this, of course, entails it is impermissible to leave B at her ragpicking status quo. If leaving B in her status quo position is impermissible and A gets B to work in his sweatshop by threatening to do so, then A coerces B, just as he does in more straightforward cases of coercion. And as in those cases, A's coercion undermines the validity of B's consent. So, in both kinds of case, B's assent is not morally valid consent because it is coerced.

In my view, an appeal to coercion is the most promising route towards explaining the wrong of most of the paradigmatic cases of wrongful exploitation. It's true that a handful of cases will involve the kind of pressure discussed in the previous section, and probably a larger number will involve fraud. But if we think there is something wrong with sweatshop employment even when workers accurately understand the terms of their contracts and are given space to deliberate about their options, then we might appeal to coercion to explain the remaining cases.

Now, note that this approach, like appeals to pressure and fraud, makes distributive unfairness morally and causally subsidiary to coercion. If, as I argued in sections 2 and 3, valid consent is morally transformative, then invalid consent is required to activate the pro tanto wrong of distributive unfairness. In this case it is coercion that invalidates consent. Furthermore, it is coercion that causally allows A to gain at B's expense; it is A's use of the impermissible threat that allows him to extract unfair terms from B.

Yet, the fact that an appeal to coercion undermines the primacy of distributive unfairness is not the only problem. A more pressing concern is that it is difficult to make the coercion charge stick in many paradigmatic exploitations. As I noted in section 2, exploitation is generally understood to be a violation of a conditional obligation, such that *if* one transacts, then one ought to do so fairly. It is commonly assumed that there is, at least typically, no *general* obligation to engage in mutually

¹¹I take it as uncontroversial that coercion undermines consent and that duress can coerce. The truth of that claim is all my argument requires. But I should mention that there is a lot of quite interesting work on *why* coercion undermines consent. See both Shaw (2012) and Dougherty (2021). In contract law, contracts involving duressed consent are typically unenforceable as unconscionable contracts. See Wertheimer (1996) and Feinberg (1986) for interesting discussions.

beneficial transactions with others. Cases involving duress are obvious exceptions. If B will die unless A saves her, it is widely, but not universally,¹² assumed A is obliged to save her. A's obligations also plausibly extend to less urgent cases as well. Working as a ragpicker might not represent an immediate threat to B's life, but its effect on her broader longevity and wellbeing might be such that it qualifies as duress.

But distributive unfairness in transactions is not coextensive with duress. If coercion is to capture the remaining cases of unfair transaction, it must be the case that when A gets B to agree to unfair terms by threatening to walk away from the transaction, his threat is impermissible because he is obliged to transact. Yet this is an implausibly demanding obligation. It cannot be the case that every unfair transaction (not covered by fraud or pressure) is also one in which A has a general obligation to transact. There will be *some* cases in which A can *permissibly* threaten to walk away from the transaction because he is under no obligation to transact. If such a threat induces B to consent to unfair terms, then (absent other defeating factors) her consent is valid and, consequently, transformative. The prevalence of such cases – of permissible exploitations – depends on our broader theories of moral obligation. But absent a very demanding view about ordinary obligations, their number is not negligible.

In fact, problems arise even for those willing to bite the bullet by arguing all paradigmatic cases of exploitation are situations in which transacting is obligatory. Consider a well-known example from the exploitation literature, Vrousalis's (2013) pit case: A finds B in a pit and can save her at no cost. If A does not save B, she will die in the pit. Suppose that in one scenario, A makes the impermissible threat to leave B in the pit unless she pays him an exploitative price for rescue. In a second scenario, A leaves B to die, making no offer to save her. Quite plausibly, leaving B to die is morally worse than coercing B into paying an exploitative price for rescue, even though the latter is also morally wrong. If the wrong of distributive fairness is activated by coercion, and this kind of coercion requires impermissible threats to not transact, then although those who use others' circumstances to exploit them *do* act wrongly, they at least refrain from carrying out the impermissible threat. However, those who don't transact at all *make good on* the threat. Like A leaving B to die, they perform a morally worse act than those who unfairly transact. We could, of course, avoid this conclusion by denying that transacting is obligatory. But this undermines the claim that the transaction involves coercion.

To recap then, fraud, coercion and pressure can all undermine the moral validity of consent. If B's assent to A's offer is predicated on deceit, vitiating pressure, or morally impermissible threats it does not constitute morally valid consent. As I noted above, when A engages in a non-consensual transaction with B he (typically) wrongs her. The problem for accounts of wrongful exploitation that ground its wrong on distributive unfairness is that many paradigmatic cases of exploitation do not involve fraud, coercion or pressure. For example, many sweatshop workers have

¹²Libertarians, for example, deny that duress, in itself, grounds an obligation to rescue. Some libertarians deny there is *any* moral obligation to rescue; others claim they are unenforceable duties of beneficence. (Of course, transaction may be required on these theories by prior contractual obligations, promises, and the like.) Consequently, for many libertarians cases involving duress are not necessarily coercive.

an accurate understanding of their contracts, have no reason to think they've been pressured into making a regretful choice (compared with their next best alternatives), and if there is no obligation to employ these workers, then threats by their employers to walk away from the negotiating table are not coercive. If consent is valid in some sweatshop cases, it is valid in many other cases of exploitation with lower stakes.

8. Valid Consent and Unfairness

Those who wish to retain the claim that distributive unfairness plays a primary role in grounding exploitation's wrong or who wish to save the intuition that many paradigmatic cases of exploitation are wrong could argue that like fraud, coercion, and pressure *unfairness* also undermines the moral validity of consent.

However, it is difficult to see how distributive unfairness – the maldistribution of benefits in a transaction – could undermine the moral validity of consent. Distributive unfairness differs fundamentally from the other factors that undermine valid consent. Pressure, deceit, impermissible threats and a lack of certain information all affect agents' deliberation, and consequently, they undermine the autonomy of choices. Distributive unfairness does not.

Consider an unfair transaction where the previously discussed factors that can undermine consent are absent. Suppose that B is a tenured academic and also happens to be a photography buff. She decides to take on a side gig working for A to earn extra income to support her expensive hobby. A offers B an unfairly low wage and B accepts. Why does B accept a low wage? Perhaps there were no better side gigs available; perhaps B found bargaining unpalatable. But whatever her reasons, let us assume none of the other factors that undermine morally valid consent are operative: B's academic position provides her a good wage so she is not duressed, she understands the terms of A's offer, A communicated the terms honestly, B's interest in photography isn't a compulsive and irresistible addiction, and so on. It is not clear how the mere maldistribution of economic gains generated by A and B's interaction undermines the autonomy of B's choice to work for A in the way that the other factors that would invalidate her consent do.

Rather than arguing that unfairness undermines the moral validity of B's consent, it seems more straightforward to claim that despite validly consenting to work for A, B's valid consent is nevertheless insufficient to override the wrong of unfair transaction. Yet, I have already argued in sections 2 and 3 that this is not so. Morally valid consent to an unfair transaction *is* morally transformative.

9. Externalities

I now turn briefly to two possible objections, beginning with the concern that my argument is essentially the same as Robert Nozick's argument against patterned principles of justice. Robert Nozick's 'Wilt Chamberlin' argument purports to show that maintaining patterned principles of justice requires intolerable restrictions on liberty, amounting to the prohibition of 'capitalist acts between consenting adults' (Nozick 1974). For Nozick, and many other libertarians, valid consent is *always*

morally transformative. My argument here is much weaker. The claim is not that *all* transactions with valid consent are morally transformed, but rather and more narrowly, that valid consent morally transforms the pro tanto wrong of unfair terms. It is possible that other pro tanto wrongs (e.g. trade in non-commodifiable goods) are not transformed by valid consent. It is also possible that they are. But in this paper my focus is squarely on unfairness. Of course, absent potentially wrong-making features like unfairness and trade in non-commodifiable goods – that is, in ‘normal transactions’ – valid consent is clearly transformative.

Further, my argument avoids the strongest responses to Nozick’s conclusion, which appeal to externalities. One reason to prohibit some exchanges, such as the sale of votes, is to limit the effect of this practice on third parties who would be harmed by the resulting inequalities of power (Anderson 1999; Scheffler 2003). We might want to limit Wilt’s profits (or tax his income) for similar reasons.

Note, however, that these arguments appeal to externalities. An equality based argument against allowing fans to pay Wilt to watch him play appeals to the effect that Wilt’s aggregate income has on relations between Wilt and others in society. The claim *is not* that by accepting a fan’s money Wilt wrongs the fan in particular. Yet this *is* what is involved in the claim that the exploited are wronged by the unfair terms of their transactions. In exploitations, the wrong of unfairness concerns the relation between the transactors themselves, not third parties.

This is not to say that exploitations cannot produce negative externalities. Of course they can. And when they do, the transactions may be all things considered wrong if the damage of the externalities is morally significant. But to reiterate, in these cases the transaction’s wrong would be grounded by its effect on third parties and not on the unfairness between the exploiter and exploited.

10. Partially Valid Consent

Finally, a second objection concerns the possibility of *partially* valid consent. Consent, plausibly, can come in degrees (Bajaj and Tomlin 2023). Suppose A threatens to flick B’s ear unless she gives him a bite of her sandwich. B really doesn’t like having her ear flicked, so she gives A a bite. Here B’s consent is, arguably, *partially* valid. Though A threatens to act impermissibly to get a bite, and thereby coerces B, we might want to say that A’s threat isn’t so significant that it completely undermines B’s consent. Plausibly, as the cost of A’s threat increases, the validity of B’s consent decreases. Consent, on this model, would not be a binary concept, but rather, a scalar one. I do not think the possibility of consent that comes in degrees undermines my arguments. There are two ways to incorporate partially valid consent.

First, we might say that if consent transforms moral permissions, and consent comes in degrees, then moral permissibility comes in degrees (Martin Peterson (2013) notably defends degrees of permissibility). If permissibility comes in degrees, then whether the wrongfulness of exploitation is transformed by consent depends on the degree of consent. Some exploitations might be more or less permissible than others. But, crucially, my argument that valid consent *modifies* moral permissions would remain. More specifically, if partial consent is incorporated this way, my first claim

can be modified as follows: the degree to which exploitation is pro tanto morally wrong is transformed by the degree to which the transactors' consent is morally valid.

Despite the strength of Peterson's arguments for degrees of permissibility, I find this view hard to accept. For this reason, I prefer the second way to incorporate partially valid consent. We can say that although valid consent comes in degrees, valid consent is transformative only when its degree passes a certain threshold. We needn't settle here where that threshold lies. What matters is that there is such a threshold. In this case, my first claim could be modified as: morally valid consent that passes a certain threshold is morally transformative, overriding distributive unfairness's pro tanto moral wrong.

11. Conclusion

This paper has defended three main claims. The first is that distributive unfairness is insufficient to ground the wrong of exploitation because morally valid consent is also morally transformative, overriding distributive unfairness's pro tanto wrong. This result is significant on its own and shows that purely distributive accounts cannot account for exploitation's wrong.

Second, I argued that distributive unfairness is not only insufficient to ground exploitation's wrong, it is causally and morally *subsidiary* to conditions that invalidate transactors' consent. Invalid consent is required to activate the wrong of distributive unfairness, making distributive unfairness morally subsidiary to these wrongs. And it is exploiters' use of these invalidating conditions that also enables them to gain at the exploited's expense, making unfairness causally subsidiary to consent invalidating conditions.

Finally, I argued that in some, and likely many, paradigmatic cases of exploitation, transactors' consent is morally valid. When transactors accurately understand the terms of their transaction, are provided space in which to deliberate, and their assent is not obtained through impermissible threats, consent is morally valid. And, consequently, even if the transaction is unfair it is not morally wrong. The scope of such cases depends primarily on the scope of our moral obligations, since coercion-enabled exploitation results from exploiters' threat to refrain from engaging in morally obligatory transactions.

Acknowledgements. I thank audiences at Warwick University, the University of Manchester, the University of Essex, and the Turkish Society for Practical Philosophy conference for helpful comments. Vuko Andrić, Sameer Bajaj, Brian Berkey, Paul Bou-Habib, Ian Carter, Sebastian Köhler, Chris Mills, Alice Obrecht, Serena Olsaretti, Fabienne Peter, Hillel Steiner, Marina Uzunova, and Mikhail Valdman also generously provided comments and suggestions on various drafts. Finally, I thank two anonymous referees for *Economics and Philosophy* for helpful suggestions that tightened and clarified my arguments.

Competing interests. There are no competing interests.

References

- Anderson E. 1999. What is the point of equality? *Ethics* **109**, 287–337.
- Arneson R. 2016. Exploitation, domination, competitive markets, and unfair division. *Southern Journal of Philosophy* **54**, 9–30.

- Bajaj S. and P. Tomlin** 2023. Consenting under coercion: the partial validity account. *Philosophical Quarterly* 74, 709–731.
- Berkey B.** 2021. Sweatshops, structural injustice, and the wrong of exploitation: why multinational corporations have positive duties to the global poor. *Journal of Business Ethics* 169, 43–56.
- Cohen G.** 1995. *Self-ownership, freedom, and equality*. Cambridge: Cambridge University Press.
- Dougherty T.** 2021. Why does duress undermine consent? *Noûs* 55, 317–333.
- Faraci D.** 2019. Wage Exploitation and the nonworseness claim: allowing the wrong, to do more good. *Business Ethics Quarterly* 29, 169–188.
- Feinberg J.** 1986. *Harm to Self*. New York: Oxford University Press.
- Ferguson B.** 2013. *The Paradox of Exploitation: A New Solution*. PhD Thesis. London School of Economics.
- Ferguson B.** 2016. The paradox of exploitation. *Erkenntnis* 81, 951–972.
- Ferguson B.** 2018. Can libertarians get away with fraud? *Economics & Philosophy* 34, 165–184.
- Ferguson B.** 2021. Are we all exploiters? *Philosophy and Phenomenological Research* 103, 535–546.
- Flanigan J.** 2018. Sweatshop regulation and workers' choices. *Journal of Business Ethics* 153, 79–94.
- Goodin R.** 1987. Exploiting a situation and exploiting a person. In *Modern Theories of Exploitation*, ed. A. Reeve, 166–197. London: Sage.
- Hart H.** 1982. *Essays on Bentham: Studies in Jurisprudence and Political Theory*. Oxford: Clarendon Press.
- Hohfeld W.** 1919. *Fundamental Legal Conceptions*. W. Cook (ed.), New Haven: Yale University Press.
- Horton J.** 2019. The exploitation problem. *Journal of Political Philosophy* 27, 469–479.
- Kates M.** 2015. The ethics of sweatshops and the limits of choice. *Business Ethics Quarterly* 25, 191–212.
- Kramer M., N. Simmonds and H. Steiner** 1998. *A Debate Over Rights*. Oxford: Oxford University Press.
- Marx K.** 1867. *Capital: A Critique of Political Economy*, volumes 1–3. Chicago: C.H. Kerr and Company.
- Malmqvist E. and A. Szigeti** 2021. Exploitation and remedial duties. *Journal of Applied Philosophy* 38, 55–72.
- McConnell T.** 2000. *Inalienable Rights: The Limits of Consent in Medicine and the Law*. New York: Oxford University Press.
- Meyers C.** 2004. Wrongful beneficence: exploitation and third world sweatshops. *Journal of Social Philosophy* 35, 319–333.
- Morishima M.** 1974. Marx in the light of modern economic theory. *Econometrica* 42, 611–632.
- Nozick R.** 1969. Coercion. In *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, ed. S. Morgenbesser, P. Suppes and M. White, 440–472. New York: St. Martin's Press.
- Nozick R.** 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
- Olsaretti S.** 1998. Freedom, force, and choice: against the rights-based definition of voluntariness. *Journal of Political Philosophy* 6, 53–78.
- Peterson M.** 2013. *The Dimensions of Consequentialism: Ethics, Equality and Risk*. Cambridge: Cambridge University Press.
- Reiff M.** 2013. *Exploitation and Economic Justice in the Liberal Capitalist State*. Oxford: Oxford University Press.
- Roemer J.** 1982. *A General Theory of Exploitation and Class*. Cambridge, MA: Harvard University Press.
- Rulli T.** 2020. Conditional obligations. *Social Theory and Practice* 46, 365–390.
- Sample R.** 2003. *Exploitation: What It is and Why It's Wrong*. Boulder: Rowman and Littlefield.
- Scheffler S.** 2003. What is egalitarianism? *Philosophy & Public Affairs* 31, 5–39.
- Shaw J.** 2012. The morality of blackmail. *Philosophy & Public Affairs* 40, 165–196.
- Sen A.** 1979. Utilitarianism and welfarism. *Journal of Philosophy* 76, 463–489.
- Sreenivasan G.** 2005. A hybrid theory of claim-rights. *Oxford Journal of Legal Studies* 25, 257–274.
- Steiner H.** 1984. A liberal theory of exploitation. *Ethics* 92, 225–241.
- Steiner H.** 1994. *An Essay on Rights*. Oxford: Blackwell.
- Steiner H.** 2010. Exploitation takes time. In *Economic Theory and Economic Thought: Essays in Honour of Ian Steedman*, ed. J. Vint et al., 20–29. London: Routledge.
- Steiner H.** 2013. Liberalism, neutrality, and exploitation. *Politics, Philosophy & Economics* 12, 335–344.
- Valdman M.** 2009. A theory of wrongful exploitation. *Philosophers' Imprint* 9, 1–14.
- Vrousalis N.** 2013. Exploitation, vulnerability, and social domination. *Philosophy and Public Affairs* 41, 131–157.
- Wertheimer A.** 1996. *Exploitation*. Princeton: Princeton University Press.

- Wertheimer A.** 2010. *Rethinking the Ethics of Clinical Research: Widening the Lens*. New York: Oxford University Press.
- Wolff J.** 1999. Marx and exploitation. *Journal of Ethics* 3, 105–120.
- Wood A.** 1995. Exploitation. *Social Philosophy and Policy* 12, 135–158.
- Yoshihara N. and R. Veneziani** 2018. The theory of exploitation as the unequal exchange of labour. *Economics and Philosophy* 34, 381–409.
- Zwolinski M.** 2007. Sweatshops, choice, and exploitation. *Business Ethics Quarterly* 17, 689–727.
- Zwolinski M.** 2009. Price gouging, non-worseness, and distributive justice. *Business Ethics Quarterly* 19, 295–306.
- Zwolinski M., B. Ferguson and A. Wertheimer** 2022. Exploitation. In *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), ed. E. N. Zalta. <https://plato.stanford.edu/archives/winter2022/entries/exploitation/>

Benjamin Ferguson is a Professor in the Department of Philosophy at the University of Warwick. URL: <https://www.benjaminferguson.org>