

Exploitation and Effective Altruism

Daniel Muñoz

Forthcoming in *Politics, Philosophy & Economics*

How could it be wrong to exploit—say, by paying sweatshop wages—if the exploited party benefits? How could it be wrong to do something gratuitously bad—like giving to a wasteful charity—if that is better than permissibly doing nothing? Joe Horton argues that these puzzles, known as the Exploitation Problem and All or Nothing Problem, have no unified answer. I propose one and pose a challenge for Horton’s take on the Exploitation Problem.

Keywords: the All or Nothing Problem, exploitation, the nonworseness claim, deontology, effective altruism, supererogation, prerogatives, moral obligation.

1. Introduction

Effective altruists: relax. You are not accused of exploiting anyone. This paper is about the moral links between exploitation and *ineffective* altruism.

To exploit someone is to take unfair advantage of them, like the sweatshop owner who pays pennies to desperate workers. The “ineffective altruist,” meanwhile, is someone who helps in a wasteful or frivolous way, like the philanthropist who picks a corrupt charity over the real deal.¹

When the consequences are bad, as they often are, it is easy to see why exploiting and wasting should be morally wrong. What’s not to loathe about the abusive sweatshop owner, or the donor who looks the other way as her cash is siphoned off by seedy middlemen?

Our feelings get mixed, however, when the acts are beneficial. Instead of the abuser, think of the sweatshop owner who gives his consenting employees a better deal than they would otherwise have (without any slime on the side—no racketeering, strong-arming local governments, etc.).

Instead of the donor funding bogus charities, suppose the gift is partly squandered but still able to

¹ To borrow a phrase from Theron Pummer (2016), the effective altruist doesn’t just care about *whether* you donate, but *where*: your donation must be reasonably effective at bringing about good results, relative to your other options for promoting the good. For a sympathetic introduction to the related movement known as Effective Altruism, see www.effectivealtruism.org; for critique, see Ashford 2018.

do a lot of good (far more than would be achieved by permissibly redoing one’s kitchen). These actions still seem wrong; we are still talking about wasted cash and grimy sweatshops. But how *could* they be wrong, if they are preferable to the permissible status quo? In the case of exploitation, this puzzle is called the “Exploitation Problem” (Horton 2019).² In the case of the ineffective altruist, it is called the “All or Nothing Problem” (Horton 2017).³

At first glance, these problems seem intimately linked. But Joe Horton, who has done more than anyone to develop them, thinks they are fundamentally separate, because there is no unified solution. His solution to the Exploitation Problem doesn’t cover the altruist, and his solution to the All or Nothing Problem doesn’t extend to the exploiter (2019: 478–79).

After laying out the problems in more detail (section 2), my first aim will be to push back against Horton’s take on the Exploitation Problem, which may only work given special bargaining conditions (section 3). We should want a general solution. My second aim is to show that there is one: a recent solution to the All or Nothing Problem can be extended to all relevant cases of exploitation (section 4). Our two problems are fundamentally one, and we can solve them in one swoop.

2. Two Problems

Horton (2019: 469) sets up the Exploitation Problem with a stylized case.

² Horton’s problem is linked to the “nonworseness claim” that a consensual, mutually beneficial action cannot be morally worse than doing nothing. This claim has been much-discussed in the literature on sweatshop labor and exploitation (Wertheimer 1996: 289–93; Zwolinski 2009; Bailey 2011; Powell & Zwolinski 2012; Arneson 2013: 393–94; Ferguson 2016; Rulli & Worsnip 2016: 213–15; Malmqvist 2017; Zwolinski & Wertheimer 2017: Section 3; Berkey 2019; Faraci 2019; Ferguson & Köhler 2019: Section 3.1).

³ The All or Nothing Problem is completely original to Horton. But cases of ineffective altruism have been around for longer (Parfit 1982; Kagan 1989: 16; Tadros 2011: 161–62; Portmore 2011: 147; Snedegar 2015: 379; Pummer 2016), along with related cases of “gratuitous harm” (Bradley 2006). For more on the All or Nothing Problem, see McMahan 2018; Sinclair 2018; Bader 2019; Lazar 2019: 242–43; Portmore 2019: section 5.2; Pummer 2019; Muñoz 2020a, Gordon-Solmon ms.

Suppose a stranger, whom I'll call Bob, desperately needs \$1,000 for a medical procedure to save his eyesight, but rather than give him your \$1,000, you take advantage of him by striking a deal: he gets the money only if he works in your sweatshop for 100 days—working “16 hours a day, seven days a week, with no breaks” (Horton 2019: 469). It seems wrong to exploit poor Bob like this. And yet, supposing he would leap at the chance to work even in a sweatshop, your offer doesn't seem worse than permissibly leaving Bob to suffer a fate he considers even more miserable.

We have a tension here. On the one hand, it appears:

Wrong to Exploit

It is wrong to give Bob the \$1,000 in exchange for sweatshop labor.

After all, by putting strings on the cash rather than giving it freely, you are taking advantage of Bob for your own profit. That said, since you didn't cause Bob's problems (and supposing you rightfully own your money), it seems:

Permissible to Keep Cash

It is permissible to keep your \$1,000.

From this, it is natural to infer that you ought to keep the money rather than exploit Bob, given the following principle, which I call:

Worse to Do Wrong

If A is wrong and B is permissible, then you ought to do B rather than A.

And yet exploitation—when it is consensual and beneficial—doesn't seem “worse” in this sense:

No Worse to Exploit

It's not the case that you ought to keep the \$1,000 rather than giving it to Bob in exchange for sweatshop labor.

This is the Exploitation Problem.

At first glance, the crux of the problem is Worse to Do Wrong. This principle links wrongness, an overall deontic property, to facts about what “ought rather” be done, which are a matter of how options fare in pairwise comparisons.

This principle is also at the heart of Horton’s puzzle about ineffective altruism: the “All or Nothing Problem.” The case is again stylized. A building is about to collapse onto two children. You could save one by holding open a small escape route, which would crush both of your arms. Or you could just as easily save both by holding open a bigger escape route (perhaps reaching a little higher), again at the very same cost of two arms.

We find a familiar tension:

Wrong to Save One

It is wrong to open the small escape route, saving only one child.

After all, you are letting the second child die for no good reason. That said, since you didn’t cause the threat to the children, it seems:

Permissible to Keep Arms

It is permissible to keep your arms, saving no one.

From this, it is natural to infer that keeping your arms is preferable to saving the one child, given:

Worse to Do Wrong

If A is wrong and B is permissible, then you ought to do B rather than A.

And yet saving one doesn’t seem worse:

No Worse to Save One

It’s not the case that you ought to keep your arms rather than saving only one life.

That is the All or Nothing Problem.

Surprisingly, even though both puzzles revolve around the same principle, Horton (2019: 478) thinks “the similarity is superficial.” Why? Because he doubts that there is a unified solution. His solution to the Exploitation Problem doesn’t generalize to All or Nothing, and his solution to All or Nothing doesn’t generalize to Exploitation.

I have two points to make. First, Horton’s solution to the Exploitation Problem seems too limited; it only applies in certain cases. Second, Horton is killing two birds with two stones; if we aim at Worse to Do Wrong, we might need only one. Let’s take these points in turn.

3. Against Horton’s Solution

Horton’s solution to the Exploitation Problem is to deny:

Permissible to Keep Cash

It is permissible to keep your \$1,000.

It is wrong to keep your money, so this is not a case where you ought to do a wrong option rather than a permissible one. Though you ought to exploit rather than keep, exploiting and keeping are both morally wrong.

But why is it wrong to keep your cash? Because, Horton says, there is a decisively better option. Bob’s 100 days of grueling labor is worth far more to you than \$1,000. So, you could spend your surplus profits (or forgo them) to make Bob’s work easier, cutting a kinder deal: Bob gets \$1,000, does less work in decent conditions, and you break even. This kind of “even deal” is fabulous for Bob and costless for you, compared to keeping the money. And Horton claims:

Costless Rescue

If you can spare someone a significant harm at no cost to yourself or others, and without violating any moral constraint, it is wrong not to do so. (Horton 2019: 473)⁴

⁴ I have replaced ‘you ought to’ with ‘it is wrong not to’.

It's wrong to keep the cash, then, because you are required to give costless benefits—and cutting an even deal counts as a costless benefit.

Horton's proposal is interesting and important. The idea is that people who are in a position to exploit, like multinational firms, are obligated at least to cut even deals with their employees—nothing less (cf. Berkey 2019). If this is right, then exploitative firms are seriously wronging the vulnerable people they contract with (because they aren't paying enough), and do-nothing firms are seriously wronging those they *do not* contract with (because they aren't paying *anything*). There exists an unconditional obligation to help others when it is possible to help them without paying a net cost.

But Horton's proposal also has a downside. As a solution to the Exploitation Problem, it may be too limited, since it only works if the exploiter *can* break even, which isn't always an option. Just imagine that, whether because of complexities of contract law or market contingencies, you have only three options: *Keep the Cash*, *Exploit Bob*, and *Give the Cash*. The same tension persists:

Wrong to Exploit

It is wrong to give Bob the \$1,000 in exchange for sweatshop labor.

Permissible to Keep Cash

It is permissible to keep your \$1,000.

Worse to Do Wrong

If A is wrong and B is permissible, then you ought to do B rather than A.

No Worse to Exploit

It's not the case that you ought to keep the \$1,000 rather than giving it to Bob in exchange for sweatshop labor.

Nor can Costless Rescue save the day. Since giving the cash is costly, it's merely optional. Horton's solution doesn't help with this version of the puzzle.

And this is not just some fringe case; it is reasonable to expect that many real-world exploiters lack the option of costless rescue. Take the case of temporary migrant labor. In Hong Kong, many foreign women toil as domestic workers for 15 hours a day, six days a week, with no hope of ever achieving permanent residence, much less citizenship—all while under constant threat of deportation: if they cannot find a new employer within two weeks of quitting a job, they must leave the country, no matter how long it has been their home (Carens 2013: 111). This is clearly an exploitative arrangement—much worse than treating migrants fairly, though not worse than closing borders. But it is far from obvious that Hong Kong’s government has the option of an even deal. What would it look like for Hong Kong to break even on its temporary worker policy? The government is dealing with uncertain tradeoffs between incommensurable goods (see Chang 2002)—economic growth vs. the benefits to immigrants. The government probably has little idea which deals would be costless. And even if such a deal were identified, it might not be politically feasible, if there is no precedent for that sort of generous arrangement. I don’t see a need to insist that, no matter the details of the case, an even deal must be on the menu somewhere, if mutually beneficial exploitation is to be wrong.

That’s my key objection. Horton’s solution to the Exploitation Problem depends on idea that the exploiter could break even. But that needn’t always be an option. The Exploitation Problem could arise even in a three-way choice between costly altruism, the status quo, and enriching exploitation.⁵ Since the possibility of cost-neutral rescue isn’t an essential feature of the exploiter’s predicament, it shouldn’t be essential to our take on the Exploitation Problem.

Now, Horton has a reply available.⁶ What if the possibility of costless rescue really *is* essential

⁵ I won’t lean on this point, but some may find Costless Rescue itself objectionable. Arguably, some cost-neutral sacrifices, like the compensated donation of bone marrow, are optional. The same may be true of certain costless favors, as when one lends or gives away something one no longer needs.

⁶ My thanks to an anonymous referee for suggesting this.

to wrongful exploitation? Horton could insist that it's not wrong to send Bob to the sweatshop in the three-option case, though it is when an even deal is possible. I think this is probably Horton's best reply. But it raises a tough question. For notice that, even in the three-option case, profiting off of Bob has the hallmarks of wrongful exploitation. You are taking unfair advantage. You are instrumentalizing Bob's economic vulnerability to enrich yourself (which suffices for exploitation, on Vrousalis's view, so long as the interaction is systematic; see Vrousalis 2013). Why *not* count this as wrongful exploitation?⁷ Horton might be able to give a persuasive answer. But until then, I think it's not unreasonable to expect that the Exploitation Problem will crop up even in the three-option case, where costless rescue isn't possible, and so we should want a solution that applies there, too.

4. Debunking “Worse to Do Wrong”

Is there a general solution to the Exploitation Problem?

There certainly *would* be, if we could debunk Worse to Do Wrong. That principle seems hard to deny. And yet, it has been denied in some recent work on the All or Nothing Problem. Let's take a closer look.⁸

The All or Nothing Problem, as a reminder, concerns a choice between (i) saving no one, (ii) saving one person at great cost to you, and (iii) saving that person plus one more, at the same cost.

The problem is that we want to say:

Wrong to Save One

It is wrong to open the small escape route, saving only one child.

⁷ A referee notes that intentions might matter here. In the three-option case, sending Bob to the sweatshop seems less bad if your goal is helping Bob (without losing money); it seems worse if your goal is making profit (no matter the toll on Bob). I agree. But Horton's reply seems to entail that even profit-hungry agents are permitted to profit off of Bob, a claim that I think needs defense.

⁸ Horton's (2017: 99) own solution is to deny Permissible to Keep Arms. If you are *willing* to lose your arms, you are *obligated* to put that sacrifice to optimal use, saving both children. As he notes, this view “does not help with the exploitation problem” (2019: 478–79).

Permissible to Keep Arms

It is permissible to keep your arms, saving no one.

No Worse to Save One

It's not the case that you ought to keep your arms rather than saving only one life.

But these are inconsistent with:

Worse to Do Wrong

If A is wrong and B is permissible, then you ought to do B rather than A.

Now, why should we believe Worse to Do Wrong? (Aren't we looking at a counterexample?)

Horton says that the principle is "intuitively correct" (2017: 96). But so are the other claims in the All or Nothing Problem. We need something more than intuitions to break the tie here. Horton suggests a second point in favor of Worse to Do Wrong, which is that:

...there are countless cases that seem to verify it. Suppose, for example, that it is permissible to say something nice, permissible to say nothing, and wrong to say something nasty. [Worse to Do Wrong] implies that you ought to say nothing rather than say something nasty. And that seems the right result. (Horton 2017: 96)

A fair point. But even false principles are *sometimes* right. We shouldn't hang on to a principle when it seems to get things wrong, unless it is supported by a deeper rationale.

Is there such a rationale for Worse to Do Wrong? Portmore thinks so:

I don't see how we can reject [Worse to Do Wrong] if we are to accept the increasingly popular view that the normative status of an action is to be explained by the reasons for and against it and its alternatives. And I, for one, don't want to reject this plausible view, as it has proved extremely fruitful in normative theorizing. (2019: 22)⁹

⁹ Portmore puts the principle a bit differently: "For any subject S and any two of her alternative options φ and ψ , if S's φ -ing is morally permissible but S's ψ -ing is morally impermissible, then S has more moral reason to φ than to ψ " (2019: 22). Ferguson and Köhler (2019: §5) also put the principle in terms of moral betterness and moral reasons (as does Ferguson 2016: 959, fn. 23). Elsewhere, I myself use "worse" as a shorthand for "less reason in favor" (Muñoz 2020a).

The idea is that we need Worse to Do Wrong to explain the grounds of wrongness. Wrongness depends on how options compare pairwise, and this dependence—Portmore tells us—somehow entails Worse to Do Wrong. But how? Portmore and others must have in mind, however implicitly and unofficially, some specific picture of wrongmaking. What is the picture?

I think it is something like this:




Fig. 1: Snedegar's Permissibility Bar

Here, Snedegar (2016: 161) illustrates the seductive idea that lies behind Worse to Do Wrong. The idea is that to be permissible is to rank high enough on the scale from worst to best (see also Fehige 1994: 47). To say that x is “better” than y , in the relevant sense, is to say x is more choiceworthy. I take this to be the same as saying that one ought to do x rather than y (in Horton’s preferred terms), or that there is more reason to do x than y (in my preferred terms).^{10,11} Snedegar’s basic idea works either way. We compare options pairwise, see which is better, and use the results to generate a ranking of options from least to most choiceworthy, which then determines right and wrong.

¹⁰ “Better” options, in this sense, needn’t have better outcomes. Pushing a man in front of a runaway trolley might result in a better outcome than letting five workers be struck, but some deontologists will insist that pushing is still the worse option. See Hurley 2017 and Portmore 2019: §1.1.5.

¹¹ Portmore’s views are more complex. He thinks that an option’s goodness correlates with the strength of the reasons to do it only when it comes to *maximal* options, as opposed to non-maximal options that can be done in different ways (2019: §6.1). This idea plays a role in his solution to a version of the All or Nothing Problem, in which not saving anyone enables you to do more good later (2019: 209–10). (When this condition doesn’t hold, Portmore thinks that it is morally better, and that you have more moral reason, to save neither kid rather than saving just one.)

At the peak of the ranking sit the best options (the ones we should and ought to choose).¹² Somewhere below is a cutoff—the “permissibility bar”—separating the wrong (gray) from the permissible (white and black). This picture entails Worse to Do Wrong straightaway. Wrong acts are below the bar; permissible acts are above; and “above” just means “better.” Hence our two problems: if doing nothing is over the bar, so is anything at least as far up—whether that be exploiting laborers, or saving only one child. Even if these acts are gratuitously bad, they are over the bar, so they can’t be wrong.

Worse to Do Wrong gets its appeal from a picture of permissible options as those “above the bar” on the scale from worst to best. If we dislodge the picture, we debunk the principle. But the picture, at least at first, seems reasonable. What’s wrong with the idea of a bar on a line?

The answer, I think, is that our best moral theories don’t reduce right and wrong to any *single* ranking: moral permissibility is multidimensional, not just “a cutoff along a single scale from naughty to nice” (Muñoz 2020). There is no one line.

The reason why you may keep your arms, even though it’s morally better to save lives, is that, in this case, “a cost to you has special significance,” not just for its moral value, but as *justification* (Muñoz 2020 and Pummer forthcoming; cf. Gert 2007 on “justifying strength” and Portmore 2011 on “moral-justifying strength”). If someone demands that you save the two kids, “It would crush my arms” is enough to show that refraining is justifiable. Now, there are different ways to say what this justification amounts to. Some would say it consists in a *non-moral reason* to self-preserve, which counts in favor of self-preservation in a non-moral way.¹³ In my view, the justification consists in a *prerogative* not to self-harm, which justifies without favoring whatsoever. For

¹² Horton uses ‘ought’ (a weak deontic necessity modal) as a synonym for ‘must’ or ‘required’ (strong deontic necessity modals). This is a common way of talking among ethicists, but there is linguistic evidence that ‘ought’ is logically weaker, as Snedegar (2016) nicely lays out.

¹³ Portmore (2011: Chapter 3) discusses this idea, but his current view also crucially uses the idea of maximal options, see his 2011: Chapter 6 and 2019. (On Portmore’s older views, see Postow 2005.)

simplicity, let us set aside non-moral reasons and stick with:

The Prerogatives Principle

You are required to do x just if, for any alternative y , the reasons to do x outweigh the combined reasons and prerogatives to do y .¹⁴

I will say that prerogatives tend to “justify” acts, or make them “easy to justify.” I call an option “justifiable” when it is made permissible by the balance of reasons and prerogatives—and “unjustifiable” when it is not. Whenever I use “reason” it will be short for “moral reason.”

But enough about terms. The real point is that permissibility depends partly on reasons and partly on prerogatives, and in general, we have a prerogative to choose less costly options. Reasons and prerogatives are two independent factors. This means that there is the potential for a mismatch—where the morally better option is favored by a weaker prerogative. When this happens, we can get counterexamples to Worse to Do Wrong.¹⁵

Take the building case. Plausibly, you have more reason to save more lives; you also have a prerogative to save zero, since that is the least costly option for you; but you have no prerogative to save one life rather than two, because saving one is just as costly. A right *not* to sacrifice doesn't entail a right to *waste* your sacrifice!

In this way, saving one can be wrong even though it's not true that you ought rather to permissibly save zero. What makes saving zero more resiliently permissible is a prerogative, and prerogatives don't make acts better or more choiceworthy.¹⁶

¹⁴ See Slote 1984, Bader 2019, and Lazar 2019 on “agent-favoring permissions,” and see Scheffler 1994, Kamm 1996, 2007, and Muñoz 2020a, 2020b on prerogatives, as well as Hurka & Shubert 2012 on prima facie permissions and Benn 2017 on supererogatory permissions.

¹⁵ Tina Rulli (2020: 377) independently came up with a version of this idea.

¹⁶ This, I think, is the only kind of case where we get counterexamples to Worse to Do Wrong—cases where A is no better than B, B is worse than C, and yet A is permissible, while B isn't, because A is protected by a bigger prerogative. Otherwise, the principle holds true. (Unless A is “on a par” with B and C, in the sense of Chang 2002; see Muñoz 2020a for discussion.)

This is a very natural way to think about the costs of altruism and how they matter; it also shows the shortcomings of the model of the “permissibility bar.” If there is just one bar, then whenever an option x is over the bar, and y is no worse, then y must also clear the bar. Every option gets the same amount of moral “leeway” before it becomes wrong; each option, in effect, enjoys the same blanket protection by prerogatives, represented by the distance between the bar and the top. This picture is hopeless if prerogatives can depend on costs. Some options cost more than others, and therefore some options (like doing nothing) can be easier to justify than others (like saving one). Whether an option is ultimately justifiable depends on the balance of moral reasons and the relative costs, not on the position of a single lowered “bar” along a single ranking. Morality is multi-dimensional, with both reasons and prerogatives, and that is why it is not always worse to do wrong.

5. Solving the Exploitation Problem

Now we are ready to solve the Exploitation Problem. We have a tension between:

Wrong to Exploit

It is wrong to give Bob the \$1,000 in exchange for sweatshop labor.

Permissible to Keep Cash

It is permissible to keep your \$1,000.

No Worse to Exploit

It's not the case that you ought to keep the \$1,000 rather than giving it to Bob in exchange for sweatshop labor.

Worse to Do Wrong

If A is wrong and B is permissible, then you ought to do B rather than A.

But now we can reject Worse to Do Wrong. We can say: you may keep the cash, you may not exploit, and yet keeping isn't better than exploiting. Again, the key is that we have two suboptimal

options, only one of which is protected by a prerogative. There is a right *not* to help, but no right to help in one way (exploiting) rather than another (gifting).

This is a simple but (I hope) decisive reply to the Exploitation Problem. There isn't pressure to accept Worse to Do Wrong if we believe in prerogatives, and so there is no conceptual tension between our three intuitions. We can coherently say that there is most (moral) reason to give Bob the cash, and least (moral) reason to keep it, even though there is a greater prerogative to keep than there is to exploit. That is how exploiting could be wrong even if it's no worse than permissibly keeping the cash.

That said, solving the Exploitation Problem is hardly the end of the philosophical problems of exploitation. We have shown that exploiting Bob *could* coherently be wrong, despite being mutually beneficial and consensual, but it is still somewhat puzzling why such an action *should* be wrong. Even if we jettison Worse to Do Wrong, there remains a tension between Wrong to Exploit and No Worse to Exploit. If exploiting Bob really is better for everyone, why should it be wrong?¹⁷

One idea is that exploitation is wrong for the same reason as ineffective altruism: it doesn't do enough to benefit others. When you send Bob to the sweatshop, you are benefiting him more than you would by not transacting. But you benefit him far *less* than you would by simply giving him the money.

True enough, but this can't be the whole story. Recall that, in the case of ineffective altruism, the agent has a prerogative to do what's in their interests, even if it's not best for others. You don't have to save the kids because it costs you your arms. Well, giving the money to Bob "costs" you the fruits of his labor. Why shouldn't you have a prerogative to *exploit* him and enrich yourself? There

¹⁷ See the sources in fn. 1 on the "nonworseness" claim. I accept that mutually beneficial exploitation can be, in some sense, consensual and best for everyone, but even when it is, I don't claim that you ought to exploit rather than keep your money; I only *deny* that you *ought to keep* rather than exploit. (Perhaps exploiting and doing nothing are on a par (Chang 2002), or perhaps you ought rather to exploit than do nothing.) My thanks to a referee for pressing me to say more here.

has to be something else going on in the case of exploitation: something wrong about it, something deeper than a mere failure of altruism, something that explains why we don't have a prerogative to enrich ourselves through exploitation.¹⁸ The problem is elusive. It isn't just that exploiting essentially involves a net harm, outright coercion, or total callousness—it doesn't.¹⁹ If we want to show that exploitation is really wrong *even when* it's beneficial and consensual, we need some deeper account of exploitation's moral core.

Exploitation, more so than vanilla ineffective altruism, remains morally puzzling. But I hope to have shown that source of the puzzle is not Worse to Do Wrong. If I'm right, we have learned something surprising by solving Horton's Exploitation Problem. What makes exploitation puzzling is not anything general about the nature of permissibility; it is something specific to the messy mix of intuitions we feel when exploiters make things better.

6. Conclusion

This paper has argued for two points about the wrong of exploitation and the obligation of effective altruism.

First, Horton's take on the Exploitation Problem may be too specific: it only applies when there is the option for an even deal, triggering an obligation of costless rescue. But some apparent cases of wrongful exploitation don't feature such an option. We may want a solution that gives a unified treatment of these cases whether or not an even deal is on the table.

Second, I've proposed a unified solution to the Exploitation Problem and the All or Nothing

¹⁸ It is widely agreed that we do not have a prerogative to enrich ourselves by violating deontological restrictions (like rights against harm); see e.g. Hurka & Shubert 2012: 9. That suggests a way to solve this puzzle of exploitation: show that exploiting, even when mutually beneficial and consensual, violates a deontological restriction.

¹⁹ If I only save one child when I could just as easily save two, I reveal that I am indifferent at best towards the second child. If I exploit Bob rather than giving him the money, I just reveal that I care *less* about Bob's experience in the sweatshop than I care about the stuff he can make for me.

Problem. You have a right not to sacrifice—but that doesn't entail a right to waste your sacrifice. You have a right not to transact—but that doesn't entail a right to exploit. Sending Bob to the sweatshop, like saving just the one child, can be wrong even though it's not true that you ought rather to permissibly do nothing. The key is that there is a moral mismatch between reasons—which in this case favor altruism—and prerogatives, which in this case protect the choice not to get involved.

There is still plenty of analysis left to do.²⁰ But thanks to philosophers like Horton, we are beginning to hear harmonies and counterpoints emerge from a cacophony of paradoxes.

“Commonsense” ethics is making ever more sense.

²⁰ I have not engaged much with the details of sweatshop exploitation (on which see e.g. Coakley and Kates 2013), nor have I said much about specific issues in the ethics of philanthropy. That said, by debunking Worse to Do Wrong, we are clearly setting the stage for a duty to give effectively rather than ineffectively.

References

- Arneson, R. (2013). Exploitation and outcome. *Politics, Philosophy, and Economics*, 12, 392–412. doi: 10.1177/1470594X13496752
- Ashford, E. (2018). Severe poverty as an unjust emergency. In P. Woodruff (Ed.), *The Ethics of Giving: Philosophers' Perspectives on Philanthropy*. Oxford: Oxford University Press (pp. 103–148).
- Bader, R. (2019). Agent-relative prerogatives and suboptimal beneficence. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics, Volume 9*. Oxford: Oxford University Press (pp. 223–250).
- Benn, C. (2017). Supererogatory spandrels. *Ethics & Politics*, 19, 269–290.
- Bailey, A. D. (2011). The nonworseness claim and the moral permissibility of better-than-permissible acts. *Philosophia*, 39, 237–250. doi: 10.1007/s11406-010-9285-2
- Berkey, B. (2019). Sweatshops, structural injustice, and the wrong of exploitation: Why multinational corporations have positive duties to the global poor. *Journal of Business Ethics*. doi: 10.1007/s10551-019-04299-1
- Bradley, B. (2016). Against satisficing consequentialism. *Utilitas*, 18, 97–108. doi: 10.1017/S0953820806001877
- Carens, J. H. (2013). *The Ethics of Immigration*. Oxford: Oxford University Press.
- Chang, R. (2002). The possibility of parity. *Ethics*, 112, 659–688. doi: 10.1086/339673
- Coakley, M. and Kates, M. (2013). The ethical and economic case for sweatshop regulation. *Journal of Business Ethics*, 117, 553–58. doi: 10.1007/s10551-012-1540-y
- Faraci, D. (2019). Wage exploitation and the nonworseness claim: allowing the wrong, to do more good. *Business Ethics Quarterly*, 29, 169-188. doi: 10.1017/beq.2018.28
- Fehige, C. (1994). The limit assumption in deontic (and prohairetic) logic. In G. Meggle and U. Wessels (Eds.), *Analyomen 1: Proceedings of the 1st Conference, Perspectives in Analytic Philosophy*. Berlin: de Gruyter (pp. 42–56).

- Ferguson, B. (2016). The paradox of exploitation. *Erkenntnis*, 81, 951–972. doi: 10.1007/s10670-015-9776-4
- Ferguson, B. and Köhler, S. (2019). Betterness of permissibility. *Philosophical Studies*. Print version forthcoming. doi: 10.1007/s11098-019-01319-8
- Gert, J. (2007). Normative strength and the balance of reasons. *Philosophical Review*, 116, 533–562. doi: 10.1215/00318108-2007-013
- Gordon-Solmon, Kerah (ms.). When saving more is worse than saving none: A study of two cases. Unpublished draft of 14 January 2021.
- Horton, J. (2017). The all or nothing problem. *Journal of Philosophy*, 114, 94–104. doi: 10.5840/jphil201711427
- (2019). The exploitation problem. *Journal of Political Philosophy*, 27, 469–79. doi: 10.1111/jopp.12181
- Hurka, T. and Shubert, E. (2012). Permissions to do less than best: a moving band. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics, Volume 2*. Oxford: Oxford University Press (pp. 1–27).
- Hurley, Paul (2017). Why consequentialism’s “compelling” idea is not. *Social Theory and Moral Practice*, 43, 29–54.
- Kamm, F. (1992). Non-consequentialism, the person as an end-in-itself, and the significance of status. *Philosophy and Public Affairs*, 21, 354–389.
- (1996). *Morality, Mortality, Volume II: Rights, Duties, and Status*. New York: Oxford University Press.
- (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press.
- Lazar, S. (2019). Accommodating options. *Pacific Philosophical Quarterly*, 100, 233–255. doi:

10.1111/papq.12252

- Little, M. and McNamara, C. (2017). For better or worse: commendatory reasons and latitude. In Mark Timmons (Ed.), *Oxford Studies in Normative Ethics, Volume 7*. Oxford: Oxford University Press (pp. 138–160).
- Malmqvist, E. (2017). Better to exploit than to neglect? International clinical research and the non-worseness claim. *Journal of Applied Philosophy*, 34, 474–88.
- McMahan, J. (2018). Doing good and doing the best. In P. Woodruff (Ed.), *The Ethics of Giving: Philosophers' Perspectives on Philanthropy*. New York: Oxford University Press (pp. 78–102).
- Muñoz, Daniel (2020a). Three paradoxes of supererogation. *Noûs*. Early online.
- (2020b). From rights to prerogatives. *Philosophy and Phenomenological Research*. Early online.
- Parfit, D. (1982). Future generations: further problems. *Philosophy and Public Affairs*, 11, 113–72.
- (2011). *On What Matters, Volume One*. Oxford: Oxford University Press.
- Portmore, D. (2008). Are moral reasons morally overriding? *Ethical Theory and Moral Practice*, 11, 369–88. doi: 10.1007/s10677-008-9110-1
- (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.
- (2019). *Opting for the Best: Oughts and Options*. Oxford: Oxford University Press.
- Postow, Betsy (2005). Supererogation again. *Journal of Value Inquiry*, 39, 245–53. doi: 10.1007/s10790-006-1344-3
- Powell, B. and Zwolinski, M. (2012). The ethical and economic case against sweatshop labor: A critical assessment. *Journal of Business Ethics*, 107, 449–472. doi: 10.1007/s10551-011-1058-8
- Pummer, T. (2016). Whether and where to give. *Philosophy and Public Affairs*, 44, 77–95. doi: 10.1111/papa.12065

- (2019). All or nothing, but if not all, next best or nothing. *Journal of Philosophy*, 116, 278–291.
doi: 10.5840/jphil2019116518
- (forthcoming). Impermissible yet praiseworthy. *Ethics*.
- Rulli, T. (2020). Conditional obligations. *Social Theory and Practice*, 46, 365–90. doi:
10.5840/soctheorpract20204189
- Rulli, T. and Worsnip, A. (2016). IIA, rationality, and the individuation of options. *Philosophical Studies*, 173, 205–221. doi: 10.1007/s11098-015-0481-6
- Scheffler, S. (1982). *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Oxford University Press.
- Sinclair, T. (2018). Are we conditionally obligated to be effective altruists? *Philosophy and Public Affairs*, 46, 36–59. doi: 10.1111/papa.12107
- Slote, M. (1984). Morality and self-other asymmetry. *Journal of Philosophy*, 81, 179–192.
- (1985). *Common-sense Morality and Consequentialism*. London: Routledge & Kegan-Paul.
- (1991). Shelly Kagan's *The Limits of Morality*. *Philosophy and Phenomenological Research*, 5, 915–917.
- Snedegar, J. (2015). Contrastivism about reasons and ought. *Philosophy Compass*, 10, 379–388. doi:
10.1111/phc3.12231
- (2016). Reasons, oughts, and requirements. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 9*. Oxford: Oxford University Press (pp. 183–211).
- Tadros, V. (2011). *The Ends of Harm: The Moral Foundations of Criminal Law*. Oxford: Oxford University Press.
- Wertheimer, A. (1996). *Exploitation*. Princeton: Princeton University Press.
- Vrousalis, Nicholas (2013). Exploitation, vulnerability, and social domination.
- Zwolinski, M. (2009). Price gouging, non-worseness, and distributive justice. *Business Ethics Quarterly*, 19, 295–306. doi: 10.5840/beq200919215

Zwolinski, M. and Wertheimer, A. (2017). Exploitation. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (summer 2017 edition). URL = <https://plato.stanford.edu/archives/sum2017/entries/exploitation/>

Acknowledgements

For helpful discussion on this paper and/or the All or Nothing Problem, I would like to thank Nathaniel Baron-Schmitt, Kieran Setiya, Chris Tucker, Theron Pummer, Kerah Gordon-Solmon, Toby Handfield, and the participants in the 2020 Winter Ethics Workshop Miniseries on the All or Nothing Problem. Special thanks to two anonymous referees for *PPE* for their clear and insightful comments, which helped me focus the paper on its core ideas; many thanks as well to Andrew Williams. My debt to Joe Horton's work is too obvious to mention, but I will mention it anyway. Finally, I am delighted to thank Mirjam Müller for invaluable comments on an early draft, and for many hours of conversation about the ethics of exploitation, which sparked my interest in the topic.

Author Biography

Daniel Muñoz teaches philosophy at Monash University, where he also helped design the new degree program in politics, philosophy, and economics. He is writing a book tentatively titled *What We Owe to Ourselves*.