

Andreas Cassee*

Trade, Exploitation, and the Problem of Unequal Opportunity Costs

<https://doi.org/10.1515/mopp-2020-0029>

Published online October 18, 2021

Abstract: This paper assesses the ‘power-induced failure of reciprocity’ account of exploitation in the domain of trade. I argue that its proponents face a dilemma. Either the cost variable of reciprocity is understood to include opportunity costs. Then, the account implausibly implies that those with more valuable outside options should get a larger part of the overall benefits of cooperation. Or the cost variable is understood to exclude opportunity costs. Then, the account has awkward implications in cases where direct costs and opportunity costs are substitutable. To evade this dilemma, the account could be amended to include a hypothetical baseline that equalizes opportunity costs. But then, the account ceases to be isolationist. Whether a cooperative interaction counts as exploitative is no longer independent of moral considerations about distributions outside the domain of trade.

Keywords: trade, exploitation, reciprocity, opportunity costs

1 Introduction

Risse and Wollner (2019)¹ propose an elaborate theoretical framework that relies on several ‘grounds of justice’, each associated with their own normative principles, to shed light on a variety of moral problems regarding trade, ranging from justice in international trade negotiations to the justifiability of corporate decisions about relocation. They draw a distinction between obligations that arise ‘in the context of trade’ and obligations that arise ‘from trading’ (11, 54). The former include all obligations agents incur while engaging in trade, including obligations that stem from other grounds of justice, such as common humanity or shared citizenship. The latter are those obligations that arise specifically on the grounds

¹ Page references throughout are to Risse and Wollner (2019), unless otherwise specified.

*Corresponding author: Andreas Cassee, Philosophisches Seminar, Universität Mannheim, L 9, 5, 68161 Mannheim, Germany, E-mail: cassee@uni-mannheim.de

that agents are involved in trade (or, to put it more precisely, on the grounds of their subjection to the global trade regime).²

The principle of justice associated with trade as its own ground of justice, Risse and Wollner argue, is a requirement to avoid exploitation. They offer a generic and a trade-specific analysis of exploitation. On their generic account, exploitation is unfairness through power. An agent is exploited in a cooperative interaction if the terms of cooperation are (a) imposed on that agent by some powerful agent, group of agents, or structure and (b) the ensuing distribution of benefits is unfair to them on some standard of fairness (88ff.). The trade-specific account builds on this generic account by specifying such a standard of fairness. Exploitation in the trade-specific sense occurs when power-induced terms of cooperation fail to be reciprocal, where reciprocity requires that the benefits of cooperation adequately reflect contributions and costs incurred by cooperators (94f.). This account applies to cooperative interactions between various kinds of agents. For example, a group of states may exploit other states by imposing a trade agreement that fails to distribute benefits reciprocally (148f.), or a company may exploit its workers by imposing wages that are too low to satisfy reciprocity (206ff.).

Risse and Wollner's theory offers an interesting middle ground between 'isolationist' and 'integrationist' approaches to trade justice. According to *isolationism*, the demands of justice in some domain (in our case, the domain of trade) are internal to that domain, that is, they are independent of moral considerations about distributions that obtain in other domains. According to *integrationism*, the demands of justice in some domain are sensitive to moral considerations about distributions that obtain in other domains.³ Risse and Wollner's approach is integrationist at the level of obligations *in the context of* trade. For example, a trade agreement may be unjust because it fails to give sufficient weight to broader concerns of poverty alleviation. But the approach is (partly⁴) isolationist at the level of obligations *from* trade. The reciprocity principle that is at the core of Risse and Wollner's 'power-induced failure of reciprocity' account of exploitation in the domain of trade is insensitive to moral considerations about the overall distribution of wealth (or any other currency of justice) among cooperators. The part of the cooperatively produced benefit each cooperator is due depends solely on the contributions they make and on the costs they incur in doing so.

² For criticism of the view that there are any obligations from trading in this sense, see Walton (2014).

³ These definitions draw on Walton's (2020) discussion of 'integrationist' and 'internal' principles of trade justice, which is inspired by Caney's (2012) earlier discussion of isolationism and integrationism in the context of climate justice.

⁴ Even though the reciprocity principle is isolationist, the account of exploitation it is embedded in need not be. Inequalities in other domains may be among the reasons why an agent has power over another agent.

One might perhaps think that the ‘mild’ isolationism of the ‘grounds of justice’ framework combines the advantages of fully isolationist and fully integrationist accounts. The framework allows us to form pro tanto moral judgements about specific interactions on a relatively thin informational basis. We need not have a full accounting of the (in)justice of the international order to call out specific transactions as exploitative. At the same time, the framework allows us to consider a wider range of issues when it comes to reaching a verdict about what justice requires in the context of trade, all things considered.

However, I will argue that the power-induced failure of reciprocity account of exploitation fails to be isolationist in a convincing way. This is because the account excludes direct moral considerations about the background distribution, while at the same time being sensitive to facts about costs and contributions that are *influenced* by the background distribution. More specifically, my objection will concern opportunity costs. Inequalities among states, companies and individuals influence the opportunity costs these agents incur when engaging in trade. If these opportunity costs are taken to be relevant for what counts as a reciprocal division of benefits, the upshot is that those with less valuable outside options have a claim to a smaller part of the overall benefits of cooperation, which seems implausible. But simply ignoring opportunity costs has awkward implications too. To save the power-induced failure of reciprocity account, we might amend it to include a hypothetical baseline that corrects for background inequality. But to do so amounts to giving up (mild) isolationism.

The remainder of this paper develops these points as follows. Section 2 discusses the principle of reciprocity that is at work in the power-induced failure of reciprocity account of exploitation and introduces the general concern that reciprocal cooperation may reproduce and exacerbate background inequality. Section 3 presents the problem of unequal outside options as one version of that general concern that is especially problematic for proponents of the power-induced failure of reciprocity account. Section 4 explores whether the problem could be avoided by excluding opportunity costs from consideration when applying the reciprocity principle. Section 5 discusses the hypothetical baseline solution. Section 6 shows how this solution is implicitly at work in Risse and Wollner’s discussion of wage justice. Section 7 concludes.

2 Reciprocity and Background Inequality

According to the power-induced failure of reciprocity account, a cooperative interaction is exploitative if and only if the terms of cooperation are (a) power-

induced and (b) fail to distribute benefits according to (or perhaps in rough approximation to) a principle of reciprocity. My focus in this paper is on the second condition.

Given its central role in their account of exploitation in the domain of trade, Risse and Wollner spend relatively little time discussing the notion of reciprocity, and they never spell out their view in the form of an explicitly stated reciprocity principle. But their remarks may suffice to reconstruct such a principle. Risse and Wollner argue that, ‘since trade is a particular type of structured cooperation for mutual gain, fairness requires all cooperation-relevant claims be satisfied proportionately’ (94). They also state that ‘cooperators have claims both in virtue of providing benefits and of incurring costs for doing so’ (94). So, the reciprocity principle will have to state a relationship of proportionality, and it has to account for costs of cooperation and contributions, as well as benefits received by each party. The question, then, is how to integrate these three variables into a statement of proportionality with two sides.

Risse and Wollner (94) briefly mention two ways in which this might be done without deciding between them. We may either rely on net benefits, subtracting the cost incurred from the benefit received by each agent, or we may add the cost incurred by each agent to the contribution they provide. Let us call the corresponding versions of the reciprocity principle the ‘net benefit formulation’ and the ‘total contribution formulation’, respectively.

On the *net benefit formulation*, reciprocity requires proportionality between net benefits and contributions:

$$b_i - c_i \propto p_i$$

where b_i stands for the benefit received by agent i , c_i denotes the cost incurred by agent i , and p_i stands for the contribution provided by agent i . In short, this version of the reciprocity principle requires that costs of cooperation are reimbursed and the remaining benefits of cooperation are distributed in proportion to contributions. If an agent provides twice the contribution of other cooperators, their net benefit should also be twice as large.

On the *total contribution formulation*, by contrast, reciprocity requires proportionality between benefits and the sum of contributions and costs:

$$b_i \propto p_i + c_i$$

On this version of the principle, there is no simple reimbursement of the costs of cooperation. Instead, costs influence the fraction of the overall benefits of cooperation each agent should receive. I take it that the net benefit formulation is more

plausible,⁵ and I will therefore rely on this formulation for my numerical examples in the following sections. But my objections apply equally to the total contribution formulation.⁶

At first sight, the reciprocity principle (on either formulation) seems to neatly separate the distribution of the benefits of a specific cooperative interaction from the overall distribution of wealth, basic goods or any other thing of value among cooperators. On closer inspection though, this separation is far from perfect. While the reciprocity principle includes no direct reference to facts about the background distribution, the facts it does rely on (facts about contributions and costs of cooperation) are very much influenced by the background distribution. This raises a concern that reciprocal cooperation may reproduce and exacerbate background inequality. Let us call this the *general problem of background inequality*.

There are two ways this problem may arise. The first and perhaps more obvious route for the background distribution to affect the distribution of cooperative benefits on the reciprocity principle is through agents' ability to provide benefits to others (p_i). An individual who could not afford college may have an inferior ability to perform some task than their fellow cooperators who enjoyed a great education. A state may have poor productive capabilities due to a small capital stock, while other states with more capital are more productive. If the distribution of cooperative benefits is to reflect the contributions provided by each party, such differences in productivity will result in differential payoffs.

How much of a problem this is, is up for debate.⁷ On the one hand, it may seem unjust that the terms of cooperation should be allowed to reproduce background inequality in this way. On the other, a proponent of the reciprocity principle might point out that a cooperator's inferior productivity, undeserved as it might be, has a direct bearing on the overall benefits of cooperation. More productive cooperators

5 Assume that A and B go fishing together. Each of them buys bait before they meet, with A paying more than B for the same amount of bait of the same quality. Each catches the same amount of fish. They then sell the fish at a local market and distribute the proceeds. According to the net benefit formulation of reciprocity, each cooperator would first be reimbursed for their respective costs, and the remainder of the money would be split equally because A and B made the same contribution. On the total contribution formulation, by contrast, A 's buying more expensive bait would entitle them to a larger share of the cooperative surplus. Intuitively, this seems unfair, and in addition, it would incentivize B to buy more expensive bait next time, thus encouraging inefficiency.

6 The problem of unequal opportunity costs that I will introduce in the next section occurs as soon as gross benefits (b_i) depend positively on costs of cooperation (c_i). This is the case on either of the two formulations.

7 As we will see in Section 6, Risse and Wollner (209f.) do think this is a problem in their discussion of wage justice, where they reject productivity wages on the grounds that an individual's productivity is often undeserved.

do not derive any *additional* benefit from the fact that their peer is less productive. If their peer were more productive, there would simply be more to distribute, and their own payoffs would remain unchanged. At best, more productive cooperators fail to *make up* for their fellow cooperator's inferior productivity, and arguably it is not their responsibility to do so.

I will not try to assess the merits of this line of argument here. Instead, I want to draw attention to a second instantiation of the problem of background inequality for which this defence of the reciprocity principle is not available. This second version of the problem relies on the influence of background inequality not on contributions (p_i), but on costs (c_i), or more precisely, on opportunity costs. As we will see, unequal opportunity costs may make it the case that the reciprocity principle allows the better off to extract *additional* benefit *because* others are worse off in terms of the background distribution.

3 The Problem of Unequal Opportunity Costs

Many of the interactions that constitute and take place within today's trade regime occur between agents with vastly unequal outside options. While some countries would do relatively well without membership in the WTO, others have little alternative to accepting the terms of membership. Some workers can gain employment elsewhere if an employer offers low wages for hazardous labour, others cannot. And while transnational corporations may decide to relocate their activities to another country if wages or taxes get too high in their current country of operation, employees generally do not enjoy similar international mobility.

There are two ways in which such facts about asymmetrical outside options may inform judgements about exploitation on the power-induced failure of reciprocity account. First, unequal outside options may make it the case that an agent has power over another agent. This is the context in which Risse and Wollner mention these asymmetries: Rich states can dictate the terms of trade because poor states have no real alternative to joining the WTO (150). Employers have the power to impose labour conditions in sweatshops because their employees have no other way out of poverty (207). Multinational corporations can threaten to relocate if employees and local communities do not accommodate their demands (223). Unequal outside options create power asymmetries, and power is one of the conditions of exploitation. On this line of reasoning, a lack of valuable outside options makes it more likely that an agent is exploited.

But there is also a second consideration that pulls in the opposite direction. Note that the net benefit an agent receives from a cooperative interaction is most straightforwardly measured as the difference between the holdings that agent has

after the interaction and the holdings the agent would have if the interaction did not take place. On this accounting, opportunity costs are a part of the costs that are deducted from the gross benefit an agent receives to calculate their net benefit.⁸ If taking part in cooperation means forgoing some other way to accrue a benefit, this counts towards the costs of cooperation. The greater an agent's opportunity costs, the greater the gross benefit that agent has to receive for the interaction to count as reciprocal and thus as non-exploitative. Conversely, the worse off an agent would be without cooperation, the smaller their share of the total benefit should be. On this line of reasoning, the fact that poor countries have little alternative to joining the WTO, or that employees have no other way out of poverty but to accept a job in a sweatshop, counts *against* the view that these agents are exploited — if their outside options are really that poor, even a small gross benefit will make for a large net benefit.

To illustrate the relevance of opportunity costs, take Risse and Wollner's (140f.) discussion of the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS). TRIPS imposes significant opportunity costs on countries with technologically less advanced economies. They forgo the opportunity to reverse engineer products that are patented elsewhere. The power-induced failure of reciprocity account may thus be employed to criticize TRIPS as exploitative. Countries with little intellectual property to protect get nothing in return for incurring this opportunity cost. But note that this opportunity cost is higher for some countries than for others. Presumably, it is middle-income countries in particular that would be able to gain technological advances by reverse-engineering, while least developed countries probably lack the technological ability to do so. If so, the account under consideration seems to imply that middle-income countries, but not least developed countries, should get a significant benefit in return for accepting TRIPS.

The exact implications of the power-induced failure of reciprocity account in this and other real-life cases depends on complex empirical assumptions. But a hypothetical example may help to bring the problem of unequal outside options into sharper focus. Assume that two agents make the same contribution ($p_1 = p_2$) to a cooperative venture with a total benefit of eight units of wealth (or well-being, or interest fulfilment, or any other currency of justice). Without cooperation, each agent would instead produce one unit in autarky (that is, opportunity costs are $c_1 = c_2 = 1$). The non-exploitative distribution of gross benefits is $b_1 = b_2 = 4$, with each agent gaining a net benefit of 3. So far, all is as it should be. But now let us assume that agent 1 gains a new outside option: They are offered participation in

⁸ One might object here that opportunity costs should be excluded from consideration when applying the reciprocity principle. This objection will be taken up in Section 4.

another non-exploitative cooperative venture, which is mutually exclusive with the cooperation under consideration, where they would gain a payoff of three units. This changes opportunity costs. While agent 2 still has an opportunity cost of 1, agent 1's opportunity cost is now 3. Non-exploitative payoffs change accordingly: Agent 1 will now receive a payoff of 5, while agent 2's payoff is reduced to 3. Note that in this case, agent 1 receives an *additional* benefit *because* there is background inequality. If agent 2 had the same new outside option as agent 1, reciprocal payoffs would be back to 4 units each.

It is not hard to see how scenarios of this kind could arise in the areas of cooperation Risse and Wollner are interested in. For example, in the context of multilateral trade negotiations, rich countries could demand a larger part of the overall benefits of trade on the grounds that a new mega-regional trade agreement among them became available as an alternative and they therefore have higher opportunity costs. Or a multinational corporation could demand wage cuts on the grounds that another country with lower wages just became available as a potential production site, thus raising their opportunity cost for keeping production facilities where they are. Note that each of these demands may be formulated as a demand for a proportional net benefit, accounting for opportunity costs — that is, as a demand for a distribution that satisfies reciprocity.

The problem here is not so much that the account under consideration yields counterintuitive judgements about exploitation in specific cases (though it may do that). The problem, rather, is how those judgements about exploitation vary with opportunity costs. If the terms of cooperation are exploitative initially, it seems implausible that the exact same terms of cooperation should be considered non-exploitative as soon as the more powerful agent gains a new outside option. Yet this is exactly what the power-induced failure of reciprocity account seems to imply.

'The less valuable options you have, the less you are given' is certainly a fitting descriptive statement about today's trade regime. But as a prescriptive statement, it is troubling, and there is no reason to believe that Risse and Wollner want to endorse it. Yet the question is whether they can plausibly avoid endorsing it, given their theory of exploitation as power-induced failure of reciprocity.

Before discussing strategies to avoid the problem of unequal opportunity costs altogether, it may be worth mentioning three mitigating considerations. Perhaps, accepting that unequal opportunity costs may in theory lead to unequal reciprocal payoffs is not quite as unpalatable as it seems, given these countervailing considerations.

First, a proponent of the power-induced failure of reciprocity account might point out that unequal opportunity costs may themselves be the result of prior exploitative trade practices. If so, perhaps we should broaden the scope of the

scheme of cooperation to which the reciprocity principle is applied. For example, if the current trade regime already unjustly favours rich states, the relevant question to ask is not whether the *additional* benefits of a *change* in the terms of trade are distributed reciprocally. Rather, we should ask whether the reformed trade regime will satisfy reciprocity overall, where this broader assessment could be either synchronic (will future net benefits be distributed in proportion to future contributions?) or diachronic (will the new regime lead to a distribution that satisfies reciprocity in the long run, accounting for past as well as future benefits, costs and contributions?). Interestingly, this line of argument may well lead us from specific questions about the merits of a particular trade agreement like TRIPS to the much more general question of how much net benefit each state receives from the trade regime overall, in comparison to a situation of autarky (perhaps correcting for how a state's productive potential in autarky is already influenced by trade).⁹

But note that past exploitative trade practices are not the only reason why opportunity costs may be unequal today. Perhaps, a state would do badly in autarky today because of reasons connected to colonialism, unjust wars, climate change, international tax competition or any other aspect of international relations that is not straightforwardly described as a part of the trade regime. If so, the trade-specific principle of non-exploitation presumably does not apply, and widening the scope of the relevant scheme of cooperation will not do the trick. Or perhaps, a state already had poor productive capabilities in autarky when it first entered international trade (if there was already a state in the relevant sense back then). If so, unequal opportunity costs today may result from a series of non-exploitative cooperative interactions that took into account unequal opportunity costs over time. Or perhaps, poor productive potential in autarky is the result of natural disaster or of bad domestic governance in the past. In any such cases where unequal opportunity costs do not result from past trade practices, the solution of broadening the scope of cooperation is not available: There is no broader scheme of cooperation to which the trade-specific principle of non-exploitation could be applied.

A second potential mitigating consideration concerns the kinds of costs that matter for judgements about exploitation. Risse and Wollner (94) stress that the costs of cooperation relevant to reciprocity include not only costs to well-being, but

⁹ Cf. Risse and Wollner's (51f.) critique of the structural equity account of trade justice defended by James (2012). According to James, the gains from trade (as measured by comparison to a situation of autarky) should be distributed equally among states, unless there are special reasons that justify an unequal distribution. Risse and Wollner's objection is in part that a state's capability to produce in autarky is already shaped by a history of interconnectedness. On the argument under consideration here, Risse and Wollner face a similar problem themselves.

also agential costs that arise when an agent's option set is diminished. Presumably, this means that we should factor in not only the next best alternative an agent gives up by cooperating (their opportunity cost in the classical sense), but the whole set of options that are no longer available to an agent because of the cooperation in question. Could these agential costs offset unequal opportunity costs?¹⁰

It is certainly possible to imagine a case where they would. For example, we might imagine an asymmetrical trade agreement where party *A* gives up a single but very valuable policy option, while party *B* loses less in economic terms, but is much more constrained in its future policy decisions by the agreement. If agential costs matter, the larger number of options given up by party *B* may offset the higher value of party *A*'s next best alternative to the trade agreement.

However, this is hardly a typical case. In most cases, the agent with the more valuable next best alternative to cooperation will likely also give up a more valuable set of options overall. At any rate, there is no reason to believe that the problem of unequal opportunity costs could be solved systematically by pointing to agential costs. As I have argued above, this problem arises because of the way judgements about exploitation vary with opportunity costs. The power-induced failure of reciprocity account allows that the same terms of cooperation count as exploitative initially, but as non-exploitative as soon as the more powerful agent gains a new outside option. The problem, in other words, is how the account judges two cases that are parallel except for changes in the opportunity costs of the more powerful agent *A*. This problem persists if we add agential costs borne by agent *B* (or any other kind of costs) to both of those parallel cases. Agential costs would only solve the problem if agent *B*'s agential costs somehow tracked agent *A*'s opportunity costs, so that *B*'s agential costs would automatically be higher as soon as *A* gained a new outside option. I see no reason why this should be the case. Therefore, while agential costs may offset unequal opportunity costs under some (rather specific) circumstances, they do not offer a systematic solution to the problem at hand.

A third mitigating consideration is that a cooperation with unequal payoffs due to unequal opportunity costs, though unproblematic in terms of reciprocity, may nevertheless be unjust on other grounds. For example, a trade agreement may be unjust not because it fails to distribute net benefits in proportion to contributions, but because it fails to give sufficient weight to human rights obligations (which are derived from other grounds of justice). The generic account of exploitation even allows us to call such an agreement exploitative, though the relevant

¹⁰ I am grateful to an anonymous reviewer for pushing me to address this possibility.

fairness principle is not reciprocity, but proportional satisfaction of human rights claims (148f.).

But once again, this line of argument will not cover all cases in which the problem of unequal opportunity costs may arise. On the grounds of justice approach, the relevant human rights claims are roughly sufficientarian in form, while exploitation in the trade-specific sense may occur above the sufficiency threshold. In such cases, the problem of unequal opportunity costs persists. Exploitative terms of cooperation may be rendered non-exploitative simply by adding a sufficiently valuable outside option to the more powerful agent's set of options.

In sum, while the implication of unequal payoffs due to unequal opportunity costs may be avoided in some instances by widening the scope of cooperation or by pointing to agential costs or to other grounds of justice, the problem of unequal opportunity costs as such remains. With this in mind, let us now turn to strategies to solve the problem more systematically.

4 Ignoring Opportunity Costs?

An obvious way to avoid the problem of unequal opportunity costs would be to simply ignore opportunity costs when applying the reciprocity principle. The costs of cooperation that matter for reciprocity, it might be argued, are restricted to direct costs. If participating in a cooperative venture implies forgoing some other way to accrue a benefit, this does not count as a cost relevant to the application of the reciprocity principle.¹¹

This move solves the original problem of unequal opportunity costs, but it has awkward implications in other regards. Note that ignoring opportunity costs is widely thought to be a failure of instrumental rationality in individual decision making. Perhaps, this need not worry us directly. Maybe reciprocity requires a different kind of accounting than self-interested rationality. But problems arise when the two kinds of accounting intersect, that is, when rational individuals who take opportunity costs into account engage in reciprocal cooperation, where reciprocity is understood to account for direct costs but not opportunity costs.

Consider the following example. Two sisters collaborate to produce vegetables on a field they own jointly. They agree that each sister will plough one half of the

¹¹ It should be noted that this solution would be at odds with Risse and Wollner's (94) insistence that agential costs matter for reciprocity, because agential costs are a kind of opportunity cost (they are about what other things an agent can no longer do because of their participation in some cooperation).

field ($p_1 = p_2$). Each sister can either rent a tractor for a cost of c_t and do their ploughing on a weekend, or they can plough by the use of their horse, though this will take them a whole week and they will have to take unpaid leave from their day jobs to do so. From the standpoint of individual rationality, the direct cost of renting the tractor and the opportunity cost of losing a week's wage are substitutable. Assuming they have no intrinsic preference for any of the tasks involved, it would normally be rational for each sister to choose whatever option is associated with the lower overall cost. Let us assume that one sister has a well-paying job ($w_1 > c_t$), so it would be rational for her to rent the tractor, while the other sister has a lower paying job ($w_2 < c_t$) and would therefore rationally choose the horse. On the version of the reciprocity principle under consideration here, the sister who uses the horse would get a rough deal, because her sister's rental cost will be split between them, while her own lost income will not.

Of course, if they know in advance that this is how costs are split, the sisters will take this into account when choosing between the tractor and the horse. Each of them will give twice as much weight to opportunity costs (which they will bear alone) than to direct costs (of which they will only bear half). The sister with the well-paying job will still rationally choose to rent the tractor. As for the sister with the lower paying job, there are two ways things could play out.

If the price for renting the tractor does not exceed twice her weekly wage ($w_2 < c_t < 2w_2$), she will rationally choose to rent the tractor. In this case, there is no unfairness in the ensuing distribution of benefits (both sisters incur the same cost and get the same benefit), but the overall costs of cooperation are higher than they could be ($2c_t > c_t + w_2$). As a result, *both* sisters will be worse off than they could be if direct costs and opportunity costs were treated symmetrically. In other words, we end up with a distribution that violates Pareto efficiency.¹²

12 Admittedly, we should not expect the avoidance of exploitation to be fully compatible with Pareto efficiency. It is an old conundrum that condemning exploitation often means condemning terms of cooperation that are mutually advantageous (e.g. Wertheimer 1996, 14). In some cases, the alternative to exploitation is no cooperation at all, and the result is a net loss to all (a violation of Pareto efficiency). But note that in such cases, Pareto efficiency is violated only because the exploiter is unwilling to engage in cooperation on non-exploitative terms. If non-exploitative cooperation did take place, it would still be Pareto efficient. In our example, by contrast, cooperation does take place, and inefficiency results not from an unwillingness to cooperate on non-exploitative terms, but from a distortion of relative prices due to the asymmetrical treatment of direct costs and opportunity costs. Therefore, there is reason to worry about the efficiency implications of ignoring opportunity costs even if one believes that the avoidance of exploitation sometimes requires us to compromise on efficiency. That being said, even if the inefficiencies entailed by ignoring opportunity costs were acceptable, this solution should still be rejected on the grounds of the unfairness problem discussed in the next paragraph.

Alternatively, if the rent for the tractor exceeds twice her weekly wage ($c_t > 2w_2$), the sister with the lower paying job will rationally opt for the horse. In this case, there is no inefficiency (both sisters choose the option with the lowest overall cost), but the ensuing distribution of benefits seems blatantly unfair. The sister with the lower paying job loses a week's wage while also carrying half the cost of the tractor rented by her sister. The sister with the higher paying job gets to keep her own wage while also being reimbursed for half her rental cost.

In sum, ignoring opportunity costs in the application of the reciprocity principle may solve the original problem of unequal opportunity costs, but the resulting mismatch between individual accounting (which takes into account opportunity costs) and reciprocity accounting (which does not) results in either inefficiency or new unfairness.

5 A Hypothetical Baseline?

If my arguments so far are convincing, proponents of the power-induced failure of reciprocity account of exploitation are confronted with a dilemma. Either they include opportunity costs in the cost variable of the reciprocity principle. Then, they encounter the problem of unequal opportunity costs. Or they ignore opportunity costs altogether. Then, the reciprocity principle has different, but equally worrying counterintuitive implications.

There is, however, a third option that may be worth exploring. Perhaps, the cost variable should be understood to include opportunity costs, but not *factual* opportunity costs that reflect the next best outside option each agent realistically has, but *hypothetical* opportunity costs that are defined with an eye to considerations of justice.

This is essentially the strategy John Roemer (1982, 1983) employs in his discussion of capitalist exploitation. Roemer introduces a hypothetical outside option to capitalist production, where (groups of) agents would withdraw from society with an equal per-capita share of alienable assets. This outside option provides a hypothetical baseline for judgements about exploitation. A group of economic agents counts as exploited if its members would be better off under this hypothetical arrangement than they are at present, and if their complements in society would be worse off than they currently are.¹³ Roemer believes that this approach is generalizable beyond the specific issue of capitalist exploitation. He writes:

13 In his 1982 article, Roemer (285) adds a third necessary condition: A group G is exploited only if its complement would be worse off if G withdrew from society with its *current* share of alienable assets.

More generally, when one says that *A* is exploiting *B*, I think one is conceiving of an *alternative arrangement* of some sort under which *A* would be worse off and *B* better off than at present. The alternative arrangement embodies one's notion of what is ethically preferable and nonexploitative ... (1982, 284f.).

Roemer's account restricts exploitation to cases where one group incurs a net loss in comparison to the hypothetical baseline, while the other group receives a positive net benefit. But a hypothetical outside option could also be introduced to amend the power-induced failure of reciprocity account. On this version of the account, an agent's net benefit from a cooperative interaction would be measured as the difference between their (actual) holdings after cooperation and the holdings they would have in some hypothetical baseline scenario, which corrects for background inequality. A cooperation would count as reciprocal if net benefits, *as measured against the hypothetical baseline*, are proportionate to contributions. For example, the two sisters' opportunity cost for ploughing with the horse would be taken into account, but these costs would be measured by applying some hypothetical (presumably equal) weekly wage, rather than their actual (unequal) weekly wages.

There are many questions about whether and how such a modified account could be made to work. For example, should the hypothetical baseline be one of complete equality, or one that only excludes morally arbitrary or only unjust inequalities? And what kinds of factors should be equalized, only external holdings such as wealth, or also internal properties such as agents' talents? I will not attempt to respond to these questions here. For our present purpose, it suffices to say that accepting the hypothetical baseline solution amounts to rejecting (mild) isolationism. What counts as a reciprocal division of the benefits of cooperation is no longer independent of moral considerations about the background distribution. Rather, such considerations inform the construction of the hypothetical baseline, and this baseline directly impacts the measurement of net benefits. In other words, to evade the dilemma regarding opportunity costs, we have to move towards integrationism.

6 Risse and Wollner on Wage Exploitation

So far, my remarks have focused on what Risse and Wollner have to say about exploitation and reciprocity in general terms. But their book also contains a discussion of the specific issue of wage exploitation (ch. 12) that both illustrates the problem of unequal opportunity costs and may be thought to provide a solution to it. Before concluding, I will therefore engage with this part of their work.

Risse and Wollner's remarks on how their account of exploitation applies to the question of wages are somewhat puzzling as they sketch two different approaches to the problem that are difficult to reconcile. At one point, they state that 'to inquire about just wages is to ask how to distribute the surplus' (208), where the surplus is defined (in accordance with textbook economics) as the sum of the differences between the company's and workers' reservation wages. The reservation wage of a company is the highest wage that a company would be willing to pay to employ a worker. A worker's reservation wage is the lowest wage that a worker would be willing to work for. These reservation wages define the shape of labour supply and demand curves. The labour supply curve is normally upward sloping (more workers are willing to work at higher wages). The labour demand curve is normally downward sloping (the company is willing to hire more workers at lower wages). The equilibrium wage is where both curves intersect. In that equilibrium, some workers will receive a wage that is higher than their reservation wage. This is the surplus gained by workers. The company will also pay some workers less than its reservation wage. This is the surplus gained by the company. The sum of these two values (the surface between labour demand and labour supply left of the point of intersection) is the total surplus. This surplus, Risse and Wollner suggest, should be distributed reciprocally.

Note that this way of framing the question of wage justice is very much susceptible to the problem of unequal opportunity costs. A worker's reservation wage depends entirely on the value of their next best alternative. A worker who has another job offer they could accept, or a piece of land they could do subsistence farming on, or some other source of income they could live off while enjoying their free time will have a higher reservation wage than a worker who does not have any of these outside options. This reservation wage is subtracted from the actual wage each worker receives to calculate their net benefit. Therefore, on the surplus accounting under consideration here, reciprocity requires that workers with more valuable outside options receive higher wages than workers with less valuable outside options.

The company's reservation wage depends on opportunity costs as well. A company will build a new production facility and employ workers there only if that promises to maximize return on capital. If an alternative investment with a higher return on capital becomes available (say, buying stock of another profitable company instead of building the new production facility), this will lower the company's reservation wage. And as the return on investment a company could receive elsewhere is priced into the company's reservation wage, this part of their profit is off the table as far as the reciprocal division of net benefits is concerned. The more profitable the outside options of a company, the lower the wages the company has to pay to satisfy a reciprocal distribution of the surplus.

In sum, on the ‘just distribution of the surplus’ approach, non-exploitative wages clearly depend on the opportunity costs of each of the parties involved. In later parts of the chapter, however, Risse and Wollner sketch a rather different approach to non-exploitative wages, which avoids this implication. They argue that wages may be exploitative in either of two ways. First, wages may be exploitative in light of being insufficient to cover the comprehensive costs of reproducing labour power (that is, the costs of living a decent life) (209). Second, reciprocity requires that the cooperative product be distributed in line with time spent producing, adjusted for training (209f.). Even though they keep talking about a fair distribution of the ‘surplus’, they no longer seem to employ the term in the technical sense introduced above. The two principles of wage justice do not seem to be concerned with the difference between market wages and reservation wages. Rather, they seem to be concerned with how the cooperative product as a whole (that is, the company’s income before subtracting wages) is distributed, and with the total wage each employee receives.

The two principles certainly make for an interesting theory of wage justice in their own right, and one with radical implications: If the cooperative product is to be distributed in line with time spent producing, it seems natural to assume that capitalist wage labour is always exploitative, as capitalists receive a part of the product with little or no time spent producing.¹⁴ But the two principles sit somewhat awkwardly within Risse and Wollner’s overall theory of exploitation as power-induced failure of reciprocity.

Officially, the costs of reproduction of labour are supposed to reflect the cost side of reciprocity, while time spent producing is supposed to reflect contributions. But both associations are far from perfect. While the reproduction of the labour force is certainly a precondition for employees’ *ability* to work, it is not strictly a cost of cooperation. The same costs of living (or living decently) arise irrespectively

14 Risse and Wollner (212) mention two reasons why capital income may nevertheless be justified on their theory of wage justice. First, capital may be thought to reflect past labour, and investors may thus have claims to some capital income on the basis of time spent producing in the past. Second, the grounds of justice framework allows for some instances of exploitation to be justified, all things considered, if they are a ‘stepping stone’ towards justice or a ‘price worth paying’. But note that both arguments are quite limited in scope. The argument from past labour will at best justify very low amounts of capital income. For example, if the founder of a company with a thousand employees invested a lifetime of their own income, this would entitle them (or those who bought stock from them) to no more capital income, spread over the entire lifespan of the company, than each of the thousand workers receive as a wage during their careers. And if the just world towards which we should strive in the long run is indeed one without capitalist labour relations, it is hard to imagine how allowing capitalist exploitation to continue could be a stepping stone towards justice.

of whether a person is employed (by a specific company, or at all).¹⁵ Time spent producing is also an awkward stand-in for contributions. Remember that the contribution element of the reciprocity principle is meant to capture the idea that the ‘provision of benefits to cooperators’ generates claims (94). Why should we say that two workers who work the same hours provide the same benefit to other cooperators, if one of them produces two pairs of shoes while the other produces three pairs?

Perhaps, then, another way to interpret the two principles of wage justice is more plausible. Perhaps, the idea that wages should suffice for a decent life is best incorporated into the reciprocity principle by way of a hypothetical outside option, rather than as a reflection of actual costs of cooperation. The theory would then say that just wages are wages workers would be willing to accept if each of them had an equal alternative way to make an income sufficient for a decent life.

Time spent producing, too, is perhaps best understood as an ‘equalizing’ measure of opportunity costs, rather than as a measure of actual contributions. If two workers spend the same hours producing, they forgo the same amount of time they could spend doing other things. On this interpretation, the time spent producing principle offers an ingenious solution to the problem of unequal opportunity costs. Rather than accepting the awkward implications of taking into account opportunity costs that are unequal, or the equally awkward implications of ignoring opportunity costs altogether, opportunity costs are measured in a way that does not allow for substantial inequality. After all, the day has 24 h for everybody.

To incorporate time spent producing into the net benefit formulation of the reciprocity principle, we presumably need a conversion factor that makes it possible to deduct opportunity costs measured in hours from benefits (wages) that are measured in monetary terms.¹⁶ This conversion factor may be thought of as a

15 A second mismatch results from the fact that the cost of reproducing labour power is meant to specify an absolute minimum wage, while the reciprocity principle is introduced as a principle of proportionality. Imagine two equally productive workers who form a worker cooperative in a situation of extreme scarcity. They spend the same hours working, and the meagre product of their cooperation is shared equally, but it is not enough to lead a decent life. Risse and Wollner (217, fn. 13) think that this would violate the principle that wages should cover the costs of reproduction of labour. If one of the workers exercised power over the other, their cooperation would count as exploitative. But if reciprocity is understood in terms of proportionality, it is not clear how this could follow: both workers receive the same net benefit for the same contribution.

16 If we use the highest attainable unitary wage as this conversion factor, the result is that the cooperation is zero sum, as measured against the hypothetical baseline. This would render contributions irrelevant for what counts as a just wage, which would make Risse and Wollner’s opposition to productivity wages compatible with the net benefit formulation of the reciprocity principle.

hypothetical equal hourly wage all individuals could gain elsewhere.¹⁷ In other words, the time spent producing principle may be thought of as a special instance of the hypothetical baseline solution.

However that may be, it is worth noting how Risse and Wollner argue for the moral relevance of time spent producing and against productivity as a determinant for just wages (which would seem to follow quite naturally from the reciprocity principle). They write:

Once we eliminate factors that are arbitrary from a moral standpoint within the context of cooperation, labor time expended is the factor most likely to survive. Unlike talent, productivity, or access to technology, time spent working generally is under individuals' control and hence a factor for which, ordinarily, individuals are responsible (209f.).

Note that many of the factors that are 'arbitrary from a moral standpoint within the context of cooperation' are arbitrary because of distributive processes *outside the domain of trade*. Undeserved differences in individual productivity may be the result of an unjust education system, racial injustice, gender injustice or unlucky familial circumstances, to name just a few possibilities. By choosing time spent producing as the determinant for just wages, Risse and Wollner want to avoid that such background inequalities influence distributive outcomes in the context of wage labour.

Risse and Wollner's account of wage justice, therefore, is integrationist, albeit in a somewhat indirect and negative way: The time spent producing principle itself does not contain any reference to moral considerations about the background distribution, but the choice of this principle is justified by such considerations. And if we take the problem of background inequality seriously in the context of wage justice (as I think we should), the question remains why we should not do the same in other contexts of cooperation.

For example, consider the case of international trade negotiations. Just like differences in individual productivity are often undeserved in the context of wages, so too are differences in country-level productivity and differences in the opportunity costs associated with joining a trade agreement. On the power-induced failure of reciprocity account, these undeserved inequalities influence non-exploitative payoffs: benefits depend on contributions (which are influenced by country-level

¹⁷ Alternatively, we might think of remuneration according to time spent producing as an implication, under special assumptions, of Roemer's (1982, 1983) hypothetical outside option of withdrawal from capitalist production with an equal per capita share of alienable assets. If labour power is homogeneous, each individual would have the same productive capability in this hypothetical scenario of autarky, so they all have equal (hypothetical) opportunity costs for an hour spent working in the capitalist economy. The conversion factor would then simply be the value of what a worker could produce in one hour in autarky.

productivity) and costs (presumably including opportunity costs). But if undeserved background inequalities should not be allowed to influence what counts as a non-exploitative wage, neither should they be allowed to influence what counts as a non-exploitative trade agreement. The argument against productivity wages Risse and Wollner endorse in their discussion of wage justice is really an argument against isolationism more generally.

7 Conclusions

In light of the enormous complexity of the economic world, it is tempting to look for moral principles that allow us to assess individual cooperative interactions independently of broader concerns about the background distribution. An obvious way to do so is to rely, among other things, on the costs agents incur when taking part in such cooperative interactions. This, however, raises the question of how to deal with opportunity costs.

While my arguments in this paper were levelled specifically against Risse and Wollner's power-induced failure of reciprocity account of exploitation, the results, I believe, are generalizable. They can be summarized by adding a third horn to our earlier dilemma. Any theorist who wants to defend an isolationist theory of trade justice that relies on costs as a basis for claims to parts of the cooperative benefit is confronted with the following trilemma: Either costs are understood to include opportunity costs. Then, the theory implies, implausibly, that those with more valuable outside options should get a larger part of the gross benefits of cooperation. Or costs are understood to exclude opportunity costs. Then, the theory has awkward implications in cases where direct costs and opportunity costs are substitutable. The third option is to rely on a hypothetical baseline that equalizes opportunity costs. While this may solve the problems associated with the first two horns of the trilemma, it means giving up isolationism. What counts as a fair distribution of cooperative benefits is no longer independent of moral considerations about the background distribution.

Acknowledgments: For helpful comments on earlier versions of this paper, I am grateful to Anna Goppel, Sabine Hohl, two anonymous reviewers, and the guest editors of this special issue of *Moral Philosophy and Politics*.

Research funding: This work was supported by Swiss National Science Foundation (100012_172668/1).

References

- Caney, S. 2012. "Just Emissions." *Philosophy & Public Affairs* 40 (4): 255–300.
- James, A. 2012. *Fairness in Practice: A Social Contract for a Global Economy*. New York: Oxford University Press.
- Risse, M., and G. Wollner. 2019. *On Trade Justice: A Philosophical Plea for a New Global Deal*. Oxford: Oxford University Press.
- Roemer, J. E. 1982. "Property Relations versus Surplus Value in Marxian Exploitation." *Philosophy & Public Affairs* 11 (4): 281–313.
- Roemer, J. E. 1983. "R. P. Wolff's Reinterpretation of Marx's Labor Theory of Value: Comment." *Philosophy & Public Affairs* 12 (1): 70–83.
- Walton, A. 2014. "Do Moral Duties Arise from Global Trade?" *Moral Philosophy and Politics* 1 (2): 249–68.
- Walton, A. 2020. "Trade Justice: An Argument for Integrationist, not Internal, Principles." *Journal of Political Philosophy* 28 (1): 51–72.
- Wertheimer, A. 1996. *Exploitation*. Princeton: Princeton University Press.