

Exploitation and Outcome

Abstract

Exploitation is interacting with another in a way that takes unfair advantage of that person. Exploitation is thought to be morally wrong even when it would bring about the best attainable outcome, hence conflicts with the consequentialist morality that holds one ought always to do whatever would bring about the best outcome. This essay aims to reconcile norms against exploitation and act consequentialism. A puzzle about exploitation is raised and resolved.

Keywords

Exploitation, unfair advantage, mutually beneficial exploitation, consequentialism, multi-level act consequentialism.

The topic of exploitation appears to be a source of embarrassment for the defender of act consequentialism.¹ Exploitation is taking unfair advantage of another person. In other words, it is using another person in a way that it unfair to that person. Doing so would seem to be morally wrong, except in extreme circumstances, as when one could save the city only by violating someone's significant moral rights. However, just as act consequentialism infamously recommends lying, stealing, cheating, betraying one's friends, and killing innocent nonthreatening persons who want to keep living, whenever doing so would bring about an outcome no worse than the outcome that would result from anything else one might instead do, so too exploitation will be morally required whenever exploiting would bring about the best reachable state of affairs.

A complication here is that exploitation is a moralized notion. To exploit is to take unfair advantage of another person, and what counts as unfair advantage? The act consequentialist has the option of maintaining that one takes unfair advantage of another just in case taking advantage of another counts as wrong according to the correct act consequentialist fundamental moral principle. On the face of it, taking up this option looks to be unpromising. Our common-sense vague ideas of treating others fairly do not run along this line.

However, exploitation is an elusive notion. No general account of what constitutes unfair advantage taking commands allegiance. A standard of exploitation would be a characterization of the phenomenon in descriptive terms, so that no further evaluation needs to be done to determine whether a given policy or action qualifies as wrongfully exploitative. If we could locate such a standard, and see that its implications for the assessment of policies and actions attract our reflective endorsement when these implications conflict with act consequentialist assessment of these policies and actions, this would be a powerful argument against the moral acceptability of act consequentialism.

We lack such a standard. In its absence a slick justification of act consequentialism might seem available. The justifying idea would be that the appeal of rules against exploitation does not challenge the claim that some version of act consequentialism is the foundational moral principle any more than the appeal of moralized injunctions such as "Do the right thing!" and "Be fair to others!" poses a challenge to act consequentialism. Lacking content, "Do the right thing!" cannot be an objection to any substantive proposed principle that claims to specify what it is to do the right thing in given circumstances. A fortiori this moralized injunction cannot be an objection to act consequentialism.

The suggested slick justification of act consequentialism fails. The problem is that many of us are disposed after reflection to adhere to the moral judgment that it would be morally wrong to extract unfairly excessive profit in dealings with others even when doing so would bring about better consequences than any alternative course of action that eschews the excessive profit-taking. Call this stable disposition to judgment *the agreed moral intuition*. To repeat the point: despite the fact that we have failed to identify a compelling descriptive standard of what it is to exploit

someone, regarding particular situations it tends to seem to us that that excessively lopsided or unequal profit-taking would be wrongfully exploitive and morally wrong even if the excessive profit taking would be here and now maximally productive of good consequences.

This essay aims to make some progress as to how to view the moral status of antiexploitation norms. A key issue turns out to be what to say about mutually beneficial exploitation. Any plausible nonconsequentialist morality will allow that people are not always obligated to do whatever would bring about the best outcome. We are often permitted to pursue our own projects even when others we could help are thereby left badly off. If we combine this idea with a moral rule against exploiting others, we are committed to the claim that given a choice between (a) having nothing at all to do with Smith or (b) engaging in mutually beneficial voluntary cooperation from which Smith gains unfairly little benefit, it follows that often it would be perfectly morally acceptable, not wrong at all, to choose course (a) but unfair and morally wrong to choose course (b). Morality then dictates that one should act in a way that would leave all affected parties worse off compared to an alternative course of action one could instead take. A candidate morality that has this implication is perverse. If all alternative candidate moralities have even worse counterintuitive implications, one might conceivably be persuaded to accept the morality that dictates (a) over (b) as the least bad alternative. But a candidate morality's having this implication is a strike against it, and a broad swath of nonconsequentialist conceptions of what we owe one another feature this implication.

The *agreed moral intuition* appears to leave us committed to a moral view that accepts the perverse implication just noted. What then should we say? Consternation looms. I shall claim that act consequentialism suggests a satisfactory resolution of this puzzle about exploitation. Here I am appealing to the familiar doctrine of multi-level act consequentialism, a doctrine that is traceable back to Henry Sidgwick and emerges clearly in work by R. M. Hare and Peter Railton (see Sidgwick 1907, Hare 1981 and Railton 1982).

Sections 3-6 of this essay survey some recent promising attempts to specify a standard of exploitation or of some significant limited type of exploitation. These attempts run into difficulties. Their failure does not show that the next proposal, or another proposal that this essay has overlooked, will not succeed. But the survey of failure does provide motivation for casting our net more widely. Section 7 shows that act consequentialism has the resources to explain the agreed intuition without exactly accepting it at face value. Act consequentialism fills the space of fundamental moral principle, and leaves no room for an antiexploitation norm, but there are levels of morality in addition to the fundamental level, and at these derivative levels rules against exploitation can have a secure place, with their content changing to some extent with shifting circumstances. Section 8 shows how the familiar doctrine of multi-level act consequentialism solves the puzzle of mutually beneficial exploitation.

1. The concept of *exploitation*.

There are two types of exploitation. In one, the exploiter gains a benefit for herself by a course of action that wrongfully harms the exploited person. In the other type, the exploiter gains a benefit for herself by a course of action that also benefits the exploited person, but the benefits gained by the exploiter are excessive by comparison to what the exploited person gains.

Exploitation can occur via an interaction between two or more people or it can occur by a unilateral imposition on the exploited. When one interacts with others, the outcome that accrues to one depends both on what one chooses to do and what the others choose to do, and this is true for each of the others as well.

A person can suffer exploitation in a transaction to which she voluntarily consents, or in a transaction not mediated by her voluntary consent.

There can also be passive exploitation, as when some receive benefits of the cooperative activity of others under conditions that trigger a duty to reciprocate, but these beneficiaries refrain from reciprocating and free ride on the cooperative behavior of the others. The free rider allows

benefits to go unreciprocated that ought to be reciprocated. This is exploitation by allowing not doing. Nor need it be the case that the appropriate context in which exploitation occurs by doing nothing must involve some prior doing by the agent that sets the stage.

2. A puzzle.

Many people accept a moral rule that forbids exploitative dealing with others and also embrace an expansive personal prerogative to pursue one's own projects when one could instead produce greater good by helping others. Such a view faces a familiar puzzle. Suppose Smith is in trouble and needs a loan, which only you can supply. You would prefer to use the money to advance your personal projects, which require an infusion of cash now. You would be willing to offer Smith a loan at what everyone would agree is an exploitatively unfair rate of interest. Ordinary common-sense morality then seems to say it would be morally wrong for you to lend Smith the money he needs at this exorbitant rate and not morally wrong for you to decline to deal with Smith at all. The problem is that Smith reasonably prefers to borrow the money from you at this rate than to have no opportunity to deal with you, and you reasonably prefer to lend to Smith at this rate than to refrain from dealing with Smith and instead pursue your personal project. Common-sense morality in these circumstances then dictates that you have a moral duty to Smith to refrain from exploiting him even though exploiting him would be mutually agreeable and would leave both of you better off, compared to the alternative scenario in which you have no dealing with Smith and pursue your own projects.

This moral posture is counterintuitive, and leads some to be suspicious of any norm against exploitation that forbids mutually agreeable trade among informed traders. Yet plausible examples of exploitative trade readily come to mind. I shall eventually claim that the multi-level consequentialist has up her sleeve an intuitively satisfying answer to the puzzle.

3. Taking advantage of vulnerability.

Allen Wood suggests that at least when one exploits weaknesses or vulnerabilities of persons, one runs afoul of a deeply rooted norm: "Proper respect for others is violated when we treat their vulnerabilities as opportunities to advance our own interests or projects." Putting this another way, he states, "It is dishonorable to use the weaknesses of others for your ends" (Wood 1995).

One can think of examples in which this formulation seems to capture what is wrong with a way of treating people. If the starving beggar falls into a ditch, one morally ought not to bargain with him, say by offering to rescue him from the ditch if he gives one his meager earnings from begging for the next year, or month, or day. But the suggested norm is too stringent to pass muster across the board. If the winters are cold where I live, I am vulnerable to whoever supplies me heating fuel, especially if there is a monopoly supplier. But a heating supply firm might charge an uncontroversially fair price for supply of a good product; there is nothing dishonorable from making a living by ameliorating the weaknesses or vulnerabilities of people if one does so fairly.² The case is clear when one's ends in the transaction are self-interested, since there is nothing dishonorable in pursuing one's self-interest within moral constraints. The case becomes blindingly clear when one's ends are reasonable and altruistic: Suppose the friars in a monastery sell heating oil to those living nearby at a reasonable price so that they can provide resources for their charitable projects that involve significantly improving the lives of distant needy strangers.

4. Fair reciprocity.

Exploitative behavior might simply be behavior that violates a norm of fair reciprocity. The two ideas seem to be paired, so we might make progress by considering what constitutes fair reciprocity.

Stuart White proposes this norm of economic fair reciprocity: “The basic idea is that where the institutions that govern economic life are sufficiently fair in terms of the opportunities they afford for productive contribution, and the awards they apportion to it, those citizens who claim the high share of the social product available to them under these institutions have an obligation to make a decent productive contribution, proportionate to their abilities, to the community in return” (White 2003). This sounds sensible. Of course we need an independent grip on the idea of what makes institutions fair, but that does not defeat White’s proposal. However, it will not pass muster if regarded as a proposed fundamental moral principle, one that holds as a matter of necessity and unconditionally at all times and places. Consider first how the principle would bear on individuals who are by nature very averse to work. Suppose I am able, but it costs me ten units of well-being to produce a small good or service working at the most efficient job that is available. Deployed so as to produce maximal value, my laboring results in net loss in value as measured by an impartial standard. In this situation, the act consequentialist will say I ought not labor, hence ought not make a decent productive contribution as required by White’s norm of fair reciprocity. In this situation requiring me to abide by fair reciprocity would be unfair.

Another case comes into view if we consider that the fellow citizens who are heaping benefits on me and to whom I owe fair reciprocity may well be less eligible recipients of my productive contributions than noncitizens. Suppose I can either work in a way that contributes to the economy of my country and benefits my fellow citizens or alternatively volunteer to aid distant needy strangers. The latter have done nothing for me, but they are badly off. Act consequentialism, especially prioritarian act consequentialism, will sometimes require that I help distant needy strangers and forego reciprocating benefits on my fellow citizens (on priority, see Parfit 2000).

A third case comes into view if we simply imagine that I am extremely badly off, one of the wretched of the earth, with wretchedness assessed in Objective List well-being terms (see Parfit 1984, 493-502). I am let us say a very poor individual in an affluent society. Its institutions satisfy the Rawlsian difference principle or some other resource-oriented justice principle that might be endorsed by an advocate of fair reciprocity, so fair reciprocity obligations are triggered. I must work to make a decent productive contribution. But in this setting act consequentialist principles dictate that I should receive and keep receiving from better off members of my society without paying them back, because even after receiving increments of their largesse I remain a morally apt recipient of further largesse. Given my constitution and circumstances, this situation endures until I die. Many versions of welfarist consequentialism will yield the result that others should do things to boost my well-being and I should also act to boost my well-being. Here welfarist consequentialism opposes fair reciprocity.

An objection that some would raise at this point is to note that I have not really argued against fair reciprocity. I have simply pointed to situations in which act consequentialism and fair reciprocity will yield opposed recommendations regarding what ought to be done. Pending some further argument, this does not add up to an argument against fair reciprocity.

The objection is correct, but there is a question about how to read the situation. My diagnosis is that the proponent of fair reciprocity is responding positively to the implications regarding policy and conduct that this norm dictates in some canonical situations. The proponent then leaps to affirmation of the norm. But in these canonical situations, many principles including versions of act consequentialism converge in their recommendations for conduct. To have a basis for making a judgment among contending rival principles, one must seek to describe situations in which they yield different and opposed recommendations for conduct. When this is done, I submit, fair reciprocity looks dubious.

A further objection holds that I am not challenging the idea that we should always treat people fairly. I am rather proposing a substitute conception of fair reciprocity that is a rival to the

one that White has in mind. Again, the objection is correct, but does not undermine my position but rather buttresses it.

Fairness in action is treating people in ways that show due concern for their interests. Fairness does not require that one be treated the same as others, but rather that one be treated as an equal, and that differences in treatment have impartial and sound justification. What principles are those? The appeal to fair reciprocity does not help. Or put this another way: The notion of fair reciprocity we should all embrace is purely formal, and says we should reciprocate whatever treatment we get from others by conforming our conduct to the correct moral principles, whatever they might be.

A subtler objection begins by challenging the background assumption (implicit in the discussion to this point) that a moral principle to be acceptable must be acceptable unconditionally and in its application to any specified set of circumstances. Maybe there are no such principles, or maybe they are epistemically beyond the ken of limited rational agents such as human beings. Be that as it may, fair reciprocity in White's formulation is not sympathetically construed as a ham-fisted attempt to concoct a principle that holds unconditionally in all possible circumstances. Instead the principle is affirmed given that some circumstances obtain. The principle mentions "citizens," so presumably one is expected to understand that the principle is intended to hold when there is a functioning state in place, so the question arises, what the citizens of that state owe to one another. Presumably other assumptions lurk in the background. One does not refute a conditional and qualified principle merely by pointing out that it is implausible if misinterpreted as an unconditional and unqualified principle.

However, the counterexamples marshaled against the White version of fair reciprocity will continue to serve as counterexamples even when we remind ourselves that the proposed norm seems to be suggested as binding only under certain conditions such as that there exists an economic system of cooperation that distributes its fruits "fairly" and perhaps also that the setting is a modern democracy. Moreover, in assessing qualified and conditional proposed deontological principles, we need to be clear exactly what antecedent conditions are being advanced that are to limit its range of application. The White fair reciprocity norm seems perfectly typical in that whatever these background assumptions might be is not made clear. A norm is being asserted, but shrouded in unstated covering. The uncertainty surrounding the content and extent of this shrouding acts to insulate the principle from objections.

5. Exploitation as opportunistic harming.

Gijs van Donselaar does not attempt to develop anything like a comprehensive norm against exploitation. But he develops a subtle account of a component of such a norm, a component he takes to be extremely plausible (van Donselaar 2009). He traces the origins of the idea to writings by John Locke, David Gauthier, and Robert Nozick (none of whom fully appreciates its implications), and uses it to criticize what he takes to be inadequate conceptions of egalitarian justice theory in the writings of egalitarian theorists (Dworkin 2000 and Van Parijs 1995). The idea is roughly exploitation as parasitism—"worsening the position of others in bettering one's own." The parasitic exploiter interacts with another in such a way that the first person benefits from interaction with the second whereas the second would be better off if the first person did not exist at all or had nothing to do with the first. The first person (the exploiter) uses the second and benefits from doing so whereas the using causes the second person to be worse off compared to the baseline of no interaction at all. For example, if one engages in an activity that dumps stuff in the river that harms someone downstream, one is not an exploiter in this sense (though there may be other grounds for regarding what one does as wrongful harming), but if one dumps harmful stuff in the river just in order to induce someone downstream to pay one to desist, and one succeeds in extracting payment for desisting, one is an exploiter in the intended sense. If one claims ownership of a hitherto unowned piece of the Earth not in order to do anything to it but simply in order to sell or lend it to someone who does want to do something

with it, and one is successful in this enterprise, one is again an exploiter on this account. If there is a queue for some scarce resource that is made open to all, and one takes a place in line just in order to sell it to someone, one again behaves in a parasitic exploitive way (one makes other people in line worse off than if one had had nothing at all to do with them). Notice that in this last case, those who are unfairly treated are not those with whom the agent has direct dealings, but third parties (who might still be involved in interaction with the agent).

One should distinguish between opportunistic agency that benefits others, serves some useful function, and opportunistic agency that just works to the advantage of the opportunistic agent. Consider what can be said in favor of ticket scalping in some circumstances. Long before a concert or sports event, tickets go on sale at set prices and the event is sold out. Suppose this market is inefficient—its outcome is not Pareto optimal, so there is scope for productive arbitrage. By buying some tickets now and selling them later to potential buyers who only later come to realize how much they want to attend the event or who always wanted to attend but are specially averse to queuing up and standing in line for tickets, one gains a profit and improves the working of this market. In realistic terms one will not be bringing about a Pareto improvement (consider the person standing in the initial line who failed to get the opportunity to purchase a ticket because the scalper was ahead of her in the queue), but one might well be bringing about a Kaldor-Hicks improvement (those who benefit from what one does could compensate the losers while still themselves gaining). Moreover, we can suppose that in the actual circumstances, there are no wealth inequality effects or other distributional considerations at work that should make us doubt that bringing about a Kaldor-Hicks improvement is really improving the situation). So in some situations, the ticket scalper is playing a socially useful role, helping to get tickets in the hands of those who value them more, and not merely harming others so that he can benefit via imposing harm on others.

Van Donsellar discusses ticket scalping, in the guise of an example involving people taking books out of the library just in order to extract payment from those who really want to use the books and are willing to pay the exploiter to return the books to the library. He does not distinguish the classes of socially beneficial and socially harmful scalpers.

The distinction between harmful using that is and is not socially beneficial shows up across the full range of types of exploitation. Consider another of van Donsellar's examples of parasitic exploitation: the person who takes advantage of a loophole in the Lockean Proviso regulating private appropriation of unowned resources, and appropriates unowned land not with the aim of using it productively but just in order to squat on it and sell to someone who does want to put the land to productive use. Again, this sort of thing might be socially beneficial even if harmful to the purchaser and arguably morally acceptable. The underlying problem here goes back to defects in the free use regime that is the background against which the Lockean Proviso operates. If a free use regime allows inefficient or depleting use of resources on unowned land provided one takes turns with other who now want to use, the predictable common result might be the ruin of the resources (Schmidtz 2006). The person who just appropriates and squats in that setting, preventing rampant depletion, is performing a useful service.

There are instances exemplifying the van Donselaar schema that look to be unfair, but mere exemplification of the schema cannot be what makes that so. There are examples of van Donselaar parasitic exploitation that do not intuitively strike us as exploitive. Consider an ordinary instance of market competition. You and I and other would be sellers of widgets are competing in the market and you win the competition. The presence of competitors for trade gives you ideas and stimulates your marketing instincts; you are better off for their presence and your interaction with them. They on the other hand would be better off if you did not exist or had nothing at all to do with them.

We might take our inclination to be tolerant of socially beneficial exploitation as a clue. Let us say that when one person harms another, a justification is required. (I would add that when one person might benefit another and does not, a justification is also required, but let that pass.)

That the act benefits third parties is a start to justification. Harming others in ways that do some good should be easier to justify than harming others in ways that do no good to anyone. On reflection, that the act benefits the agent herself is also a start to justification (better that I benefit than that no one benefits).

The ultimate justification of an act that harms another is that compared to the available alternative courses of action this act brings about an overall outcome no worse (as assessed from an appropriately impartial perspective) than the outcome that would result from anything else one might instead have done. So anyway holds the act consequentialist.

Another type of example involves assignment of property rights in land to disabled persons. They have no interest in working the land but can sell or rent it to someone who does want to use it productively. Without the existence of the disabled, the productive would simply divide the land among themselves and do better than they would if the disabled did not exist at all, whereas the disabled quite obviously are better off due to the existence and interaction with the productive. One man's parasitism is another man's distributive justice. Interacting with me in a way that leaves me better off and you worse off than you would have been had you had nothing at all to do with me or if I did not exist does not signal that I am conducting myself in an unfair way, much less constitute my conducting myself in an unfair way, if I have a right to exist and you are under some obligation to better my condition.

Whether interacting with others in a way that makes the agent better off via worsening the position of others relative to the baseline in which the agent does not exist or they have nothing to do with him is morally wrong depends on further features of the situation. For example, the others may be obligated to interact with the agent on these terms. Consider

THE DIFFICULT, TROUBLED PEOPLE

Some people possess psychological traits that amount to rational agency capacity at a level that undeniably renders them persons, full members of the moral community, along with other, unfortunate traits. They are poor transformers of resources into goods for themselves and for others. They are poor learners of more productive ways of interacting with others. Their lives are not very happy and fulfilling. Nonetheless they ought not to be ignored or neglected. Others ought to interact with them, even though generally the difficult people will gain from these interactions, will benefit from the existence of those with whom they interact, and those interaction partners would be better off if these difficult people did not exist or had nothing at all to do with them.

There are act consequentialist theories that will readily endorse what van Donsellar is forced to identify as exploitative parasitism on the part of the difficult, troubled people, in a wide range of likely circumstances. For example, consider prioritarian theories that hold that what is morally right is maximizing a function of aggregate human well being that assigns greater gains to gaining a benefit for the person, the worse of the person would otherwise be in lifetime well-being. Yielding this implication for the question how we ought to treat difficult troubled people is a large plus for those versions of act consequentialism. In this sort of situation the fact that one's act amounts to van Donsellar exploitation is not a wrong-making feature of it.

One might stiff-arm this last assertion. One might claim that it is always pro tanto morally wrong to engage in van Donsellar parasitic exploitation, but countervailing reasons can overcome this moral presumption. The act consequentialist will claim that the countervailing reasons will involve the balance of consequences and the balance of consequences always trumps, leaving the pro tanto consideration with no work to do.

The van Donsellar idea is connected to a broader and perhaps less controversial idea of harmful unfair using. Consider Warren Quinn's idea that there is a special moral presumption against harmful opportunistic agency. Quinn had distinguished first between "agency in which harm comes to some victims, at least in part, from the agent's deliberately involving them in something in order to further his purpose precisely by their being so involved (agency in which they figure as *intentional objects*) and harmful agency in which either nothing is in that way

intended for the victims or what is so intended does not contribute to their harm.” He further distinguished between “direct agency that benefits from the presence of the victim (direct *opportunistic* agency)” and “direct agency that aims to remove an obstacle or difficulty that the victim presents (direct *eliminative* agency)” (Quinn 1993, at 184 and 186). According to Quinn, harming by direct agency is harder to justify than harming by indirect agency and harming by direct opportunistic agency is harder to justify than harming by direct eliminative agency. A qualification: According to Quinn, the presumptions against direct agency and opportunistic direct agency come into play only when the agent harms another and thereby violates a moral right possessed by the harmed person. These presumptions do not generate exceptionless restrictions but are simply moral considerations (pro tanto reasons) to be balanced against other moral considerations in determining what all things considered is morally permissible and morally required to do.

This is a plausible set of ideas. It might be correct. However, notice three points. The Quinn proposal can be applied only in conjunction with a prior specification of people’s moral rights. Important moral work is done off-stage. So the idea does not amount to a descriptive characterization of wrongful exploitation that could serve as a standard of the sort I claim is not to be found. Second, Quinn is seeking a wrong-making characteristic, not a determinate principle that holds always and without exception, and one might worry he is setting a low bar for success. If we end up with a bunch of pro tanto considerations and no generalizations organizing and constraining their moral force in determining verdicts as to what to do in particular situations, what gives us confidence that the pro tanto considerations are not just prejudices or cultural givens? This question is explored in section 7 of this essay. A third point is that one might find it puzzling that aiming to involve people in one’s plans to one’s benefit in ways the harm them is worse than just harming them. One possibility is that we are specially tempted to use people for our benefit in ways that are generally overall harmful, so a rule that bends the twig by establishing a strong presumption against such using might be helpful overall in promoting good consequences. If this should turn out to be the best rationale for the rule, it would not then challenge but rather illustrate consequentialist moral thinking.

6. Exploitation as unfair division of benefits

In a thorough discussion of how one might identify a moral principle that determines, when people voluntarily interact in mutually beneficial ways, whether the distribution of benefits and burdens is excessively lopsided so that the interaction should count as unfair and exploitive, Alan Wertheimer tentatively embraces the idea that the hypothetical terms of exchange that would obtain in a perfectly competitive market for the good or service in question provide the standard for a nonexploitive exchange. In this connection he notes first that we should distinguish between unfairness in the background conditions that lead to transactions and unfairness in the transactions themselves. He observes further that “even though a competitive market price does not reflect a deep principle of justice, it does reflect a crucial moral dimension of the relationship between the parties to the transaction. The competitive market price is a price at which neither party takes *special* unfair advantage of particular defects in the other party’s decision-making capacity or special vulnerabilities in the other party’s situation. It is a price at which the specific parties to the particular transaction do not receive greater value than they would receive if they did not encounter each other. It may or may not be a “just price,” all things considered, but it may well be a nonexploitative price, for neither party takes unfair advantage of the other party” (Wertheimer 1996, 232). Wertheimer adds that when parties A and B engage in exchange, “in a competitive market, A does not have the space to transact with B at a different price, even if A would prefer to do so.”

Wertheimer first of all makes a mistake about the idea of a competitive market. It does not force compliance with market price as Wertheimer supposes. In a competitive market, nothing forces me to insist on the competitive price for a good or service I am selling, and

perhaps it would be unfair of me not to sell below the market price. Of course, if I sell for less, the market is no longer fully competitive, but this just illustrates that the going market price does not twist my arm. The same is true for buying. The competitive price does not block me from paying more. There is a competitive market for housecleaning in San Diego, with many buyers and many sellers, but nothing blocks me from my paying the person who cleans my house some multiple of the going rate. And what holds here holds generally.

Wertheimer writes, “McDonald’s might want to pay its workers a higher wage, but it cannot afford to do so if it would thereby lose all its customers to Burger King.” Let us assume that if the market is currently competitive, if McDonald’s charges more for its products, it will lose all its customers to competitors. But one can raise the wages of the unskilled workers one employs without raising prices. One can offset rising labor costs by lowering executive pay, or the pay of skilled workers and managers, or dividends paid to shareholders, or some mix of these policies. In a truly competitive market, of course, one could not enact any of these policies, but “market forces” do not constrain us from acting in ways that would render the market no longer perfectly competitive, and perhaps we have moral reason to do so. If one can behave in this way, any number of people can. For that matter, nothing prevents would-be consumers of a good or service from paying a higher price for it just because they see that the higher price is due to the giving of a higher than necessary payment to suppliers or laborers and acknowledge that this payment is morally mandatory. Even if they do not behave this way, maybe they should.

The idea of a perfectly competitive market assumes that purchasers of goods seek to maximize the benefits they gain from their purchases and sellers seek to maximize their profit. These assumptions about the character of agents’ preferences are just empirical assumptions justified by their value in predicting agents’ behavior and their overall economic effects. My point is simply that the assumptions a social scientist makes about your preferences do not in any way constrain your choice about how to behave when confronted with a market setting in which there are many buyers and many sellers of each good that is traded and prices are competitively formed.

Of course, these specific empirical assumptions about the shape of people’s preferences can be relaxed, and would need to be if many individuals’ preferences shift. If some consumers become willing to purchase running shoes made in China only if they are produced by a manufacturing process that is safe for their workers, from the economists’ point of view running shoes made safely are now a unique good that will fetch its own price on the market. The same would be true if some consumers developed racist preferences that influenced their consumer behavior, of course. The market neutrally registers the preferences of economic agents. This drives home the point that there is nothing normative about market prices and there should be no presumption that competitive market prices are by any substantive moral standards fair or nonexploitative.

Suppose it is claimed that a firm that sells a product to consumers at the competitive price takes no special unfair advantage of its consumers. Is there anything to this thought? It is uncontroversially true that if market exchange is perfectly competitive, no agent has even a trace of monopoly power, so none can engage in successful hard bargaining exploiting special bargaining leverage to gain extra profit. If I try to pay my workers less, they will immediately quit and find jobs exactly as good as the status quo ante at my firm. But these facts confer no presumption of fairness on perfectly competitive market prices.

To illustrate the point, consider the market for housecleaning services in a border city. There are many buyers and many sellers; the market is reasonably competitive. The background facts that set the contours of equilibrium process in this market include the facts that many people on one side of the border are born in poverty and lack access to good education and good jobs and many people on the other side of the border are by comparison enormously lucky and wealthy. As a comparatively lucky and wealthy home owner, I am not responsible for these background facts that set market prices, but why is it morally acceptable to benefit from them by taking the price as

given and paying just the going rate? Suppose my reservation price for the service is far above the market rate. What renders it morally acceptable, just and fair, for me to gain a lot of consumer surplus when I purchase housecleaning services? The same question reverberates across the board whenever one purchases goods and services in a global marketplace (see Cohen 2000, chapters 9 and 10).

(In passing, we notice that in this conversation a nagging act consequentialist question can be heard: rather than fix attention on the particular persons with whom one is engaged in transactions, why not drive whatever hard bargains one can and then use the profits in whatever way is maximally beneficial? Why not drive a hard bargain with my moderately poor housecleaner (or clean my own house) and send the money thereby saved to the charity that will most effectively use it to achieve real gains for people elsewhere who are far worse off?)

One might claim that a perfectly competitive market will yield outcomes that have the morally attractive characteristic of Pareto efficiency: at the equilibrium of market trading one cannot alter the outcome to make anyone better off without making someone else worse off. But first of all, efficiency in this sense, even if necessary for fairness, is surely not sufficient for it. And a second point is also relevant: The conditions that qualify a market as efficient in the Pareto sense are more stringent than the conditions that qualify a market as perfectly competitive. The advocate of the perfectly competitive market as containing a standard establishes that transactions are nonexploitative cannot wrap herself in the thin shawl of Pareto efficiency.

7. Consequentialism and norms against exploitation.

These various suggested specific substantive conceptions of fairness that might give specific content to a norm against exploitation are a motley. They take up different domains of social interaction at different levels of abstraction. One might suspect that nothing much follows from the conclusion that this collection of proposals is to be rejected. Maybe the next proposal will prove adequate; or maybe some other proposal already advanced but ignored in this discussion is adequate. Maybe so. However, there is another way to view the argumentative situation. Attempts to develop a substantive conception of what it is to behave exploitatively, one with content for determination of what is right and wrong, peter out unsuccessfully after examination. This conclusion echoes the conclusions of authors that are firmly in the deontological camp and not pressing a consequentialist ideology. Wertheimer's book is the most thorough recent treatment of the topic and by his own admission he comes up with no results on the substantive content issue (at least if we ignore his tentative and anyway blatantly unsuccessful flirtation with a hypothetical competitive market price norm).

At this point the argument in this essay faces a challenge. The various proposed conceptions of exploitation have been rejected as inadequate on the ground that there are circumstances in which all things considered it is implausible to regard violating the proposed norm as morally wrong. The recipe employed to generate the counterexamples is to find cases in which adhering to the proposed norm would expectably lead to bad outcomes. This line of argument invites the response that when sympathetically interpreted, the criticized antiexploitation norms will not purport to be universal exceptionless moral principles but rather pro tanto considerations relevant to determining correct moral choice. So interpreted, the suggested conceptions are not vulnerable to my counterexamples, and for all that has been said, might all be true. Moreover, they stand as strong reasons to reject act consequentialism. Even if the pro tanto norms against exploitation should sometimes give way when the consequences of adhering to them are very bad, they surely have some weight against consequentialist calculation, and this suffices to rebut the act consequentialist skepticism about antiexploitation norms (or about any other deontological norms, for that matter).

In reply: I have urged that when the suggested antiexploitation norms do not coincide in their implications for conduct and policy with the implications of reasonable versions of act consequentialism, they look implausible. At the level of casuistry, I deny what the critic of act

consequentialism asserts. But this looks to be a standoff of opposed intuitions, not a clear-cut victory for the act consequentialist.

A second reply: morality as a collection of pro tanto considerations is implausible. If all that one can say, in defense of the claim that it would be wrong to do act X in this particular situation, is that one's uncodifiable judgment is that the overall balance of pro tanto reasons rules out doing X, any claim one might make to the effect that one's judgment is objectively correct looks arbitrary and indefensible. One's claim that a moral judgment about a particular choice of action is correct is vindicated only if one can show that the judgment is derivable from a full specification of the relevant empirical facts plus moral principles that state universal general moral claims that are affirmed in reflective equilibrium after full scrutiny and deliberation. To be suited for this role as moral principle, a generalization must specify what determines permissible and required action in descriptive terms that do not need further uncodified moral intuition to be employed. The relevant empirical facts that need to be specified in determining what to do are those picked out by the principles under review as relevant. One can object to candidate moral principles on the ground that their implications for conduct in particular circumstances are counterintuitive. But these intuitive objections are ultimately grounded by discovering moral principles that explain and justify them. If a search for principled defense of intuitions one accepts turns up empty pockets, the intuitions begin to lose their force as objections to candidate principles. This essay has argued that accounts of antiexploitation norms that aim to give systematic order to our intuitions concerning particular cases of unfair and exploitative transactions turn up empty pockets.

The views stated just above are controversial along many fronts, and cannot be properly elucidated, much less defended, in this essay. (For contrary views, see Ross 1988 and Dancy 2004). For purposes of this essay, the claim advanced is hypothetical: if you accept the claims in the previous paragraph, you will not find intuitive objections to candidate moral principles—be they consequentialist or nonconsequentialist—on the ground that they do not chime in with our judgments about the wrongness of particular types of exploitative treatment of people to be capable of defeating the candidate principles unless they lead to a more intuitively acceptable candidate set of principles.

However, this line of thought is too swift a dismissal of the challenge our common-sense ideas about exploitation pose for the consequentialist. What I have called the “moralized generic” norm against exploitation is not entirely lacking in content. The norm affirms an agent-relative moral duty to refrain from (1) harming another in order to gain profit for oneself unfairly and (2) interacting with another person in a way that generates for oneself an unfairly large share of the gains from interaction. In particular situations we do have strong inclinations to judge that particular actions that harm others and benefit us are unfair and exploitative and that some actions channel top ourselves an unfair share of the cooperative gains. Even in the absence of a determinate descriptive standard specifying what constitutes exploitation, these intuitive judgments pull against adherence to any version of act consequentialism.

The arguments offered in this essay against proposed standards of exploitation tend to be counterexamples that appeal to maximizing consequentialist judgments. This perhaps gives a hint about how best to proceed. We explain away the moralized antiexploitation intuition by interpreting it as indirectly derived from consequentialist principles. In this view the intuition is not strictly speaking correct, but we can understand why we are prone to accept it, and in a roundabout way the intuition turns out to be acceptable in that we should not attempt to expunge it from our moral consciousness. The act consequentialist is here following a soft debunking strategy. Tipping the argumentative balance, the act consequentialist can offer a further consideration: she has a satisfying explanation of the puzzle described above in section 2.

The starting point is a very general fact. Growing up in any of a wide array of human cultures, we would be trained into acceptance of broadly deontological moral views. These views posit moral rules that condemn specified types of action and hold that it is morally wrong to do an

act of this type and that the mere fact that doing so in some situation would lead to a somewhat better outcome than would otherwise obtain does not suffice to render the act permissible. A moral rule against exploitation or against violation of fair reciprocity often figures in the set of rules inculcation of which is part of normal upbringing in the culture. We should not romanticize these rules. As John Stuart Mill notes, “Wherever there is an ascendant class, a large portion of the morality of the country emanates from its class interests and its feelings of class superiority” (Mill 1978 [1859], at 6). The rules in any given setting might also reflect the contingent outcome of social struggles. But nor should we demonize the rules. We should expect them to some degree to incorporate the local sense of what we truly owe one another, achieved by distancing ourselves in thought from distorting factors. What is upheld on the form of moral rules for individual conduct is reflected to some extent in informal social norms and legal codes enforced by governments.

From an act consequentialist perspective, these entrenched rules reveal significant merit. Since human persons tend to be selfishly partial, not well informed about matters of fact that bear on proper choice of conduct, and anyway not very good at sensibly integrating such pertinent empirical facts of which they are aware into the determination of what to do, promulgating a social morality for guidance of individual decision making that consisted in the single rule “Do whatever would bring about the morally best outcome” would lead to bad choices and bad outcomes. On the whole better choices are made and better outcomes achieved if people generally guide their individual decision making by the going moral rules and by the associated social norms and legal code. “Follow the established rules” is generally a good decision guide even if the established rules are very imperfect as assessed from an act consequentialist perspective. So someone who wishes to be successful in living her life so as to produce best consequences is well advised to follow the going rules, and acts of establishing such rules and of sustaining them by conforming to their requirements tend to be better by the act consequentialist standard for judging actions that other decision guides one might instead follow. (The degree to which this is so in a society depends of course on the quality of the specific rules in force.)

To see how agent-relative rules might prove to be fairly good act-consequentialist guides, consider a humdrum example of self-interested action that can be justified in act consequentialist terms but falls in a class of closely similar behavior that cannot be. Suppose highway traffic is heavy; the lane in which I am traveling is temporarily stalled, and I veer into the adjoining lane, inducing the driver behind me to apply his brakes and intending to benefit myself by the combination of my lane change and his slowing down to allow me to move faster through the congestion without causing a car accident. No doubt many such actions are selfish and morally wrong; all things considered they impose more cost (especially car collision costs) on others than gain to the selfish driver. But in some humdrum cases I have especially strong reason to get through the traffic fast and the gains to me from acting selfishly outweigh and are reasonably expected to outweigh losses to all possibly affected others. (If I have special reason to suppose I am headed toward very low lifetime well-being compared to those whose well-being I will be slightly diminishing, the straight aggregative utilitarian case for my selfish behavior is increased.) The thought that there should be a norm against gouging small advantages for oneself when driving in congested traffic sits easily with the further thought that sometimes such action is, in the circumstances, the choice that leads to the best available outcome and hence morally justified.

The idea of one act being better than another by the act consequentialist criterion of morally right action requires a comment. One might suppose that act consequentialism just identifies, of all the disjoint acts an individual might do at a time, which one would bring about the best outcome (according to the standard of assessment of consequences that the principle incorporates. Unless more than one act is tied for best by this standard, act consequentialism seems to just identify one best act. So how do we get to comparisons of what is done as better or worse? Recall that in his canonical formulation of a version of act consequentialism Mill states that acts “are right in proportion as they tend to promote happiness; wrong as they tend to

produce the reverse of happiness” (Mill 2001 [1861], at 7). One way to read Mill’s idea here is that he is suggesting that an act is “righter” or “wronger,” depending on the shortfall between the value of the outcome it would produce and the value of the best outcome the individual could instead have brought about by choice of some alternative action. The greater the shortfall, the more wrong the act.

Even in a society whose socially established moral rules, social norms, and legal code were all ideal by a consequentialist measure (meaning that establishing and upholding those very rules, norms, and code in the actual circumstances faced would bring about the best available outcome), there would very likely be large differences between the content of this ideal set and the act consequentialist principle that we are supposing is the sole fundamental moral principle that determines what is morally right and wrong to do. What we ultimately have most reason to do in any situation in which we might find ourselves is that act that is picked out as best and hence morally right by the act consequentialist standard. But following established rules in this imagined society, and for that matter in the actual society we inhabit, is a generally good strategy for picking out acts that are better rather than worse (more rather than less right) by the act consequentialist standard.

The suggestion then is that a generic norm against exploitation, using others unfairly, is very likely to be a stable component of a set of moral rules socially established across many cultural settings, that is approved as part of multi-level act consequentialism. It is a component of the actual set we are likely to have and would also be a component of that set if it were reformed to better serve act consequentialist ends. Moreover, specific substantive conceptions of exploitation, even though they would not pass muster if proposed as candidate elements in a set of fundamental moral principles (that ultimately determine what morally ought to be done, what is genuinely normative for action), including some canvassed in this essay, may well make good sense in particular contexts, in a given society at a given time. The various substantive anti-exploitation norms would make sense in that establishing and upholding them would tend to have reasonably good consequences as compared to feasible alternatives.

Accepting that there can be a consistent multi-level doctrine that encompasses a fundamental act consequentialist principle and deontological nonconsequentialist social norms and legal codes, how do we avoid inconsistency if we endorse both act consequentialism and a morality of rules as our moral commitments? Either exploiting a person is morally wrong or it is not. Act consequentialism says that the fact that the act you might perform qualifies as exploitative has nothing to do with whether the act is morally right or wrong—only the consequences that would ensue fix that matter. The morality of moral rules including rules against exploitation on the contrary says that the fact that the act in question qualifies as exploitative establishes that the act is morally wrong or at least pro tanto morally wrong. We cannot hold both that exploitation status is and is not a wrong-making factor.

Consistency is regained if we note that one can uphold a nonexploitation norm as part of the public morality that is officially proclaimed and also maintain that at the fundamental moral level nothing other than the act consequentialist principles has a place. “X is impermissible according to public morality” is consistent with “X is morally permissible according to fundamental moral principle.”

Moreover, strictly speaking, a moralized generic norm against exploitation is contentless. It may serve as a reminder and pointer that actually guides action and guides it well. But if one holds that it is wrong to interact with a person in a way that is unfair to her and imposes no substantive constraint on what is to count as treating someone unfairly then there is no conflict between what one holds and upholding the moral judgment that is dictated in each and every case by act consequentialist principle. Upholding the combination brings it about that at the fundamental level, interacting with someone in a fair way is behaving toward her only as act consequentialist principle dictates.

Is this wildly implausible? Keep in mind that the claim is not the denial that well brought up people in tolerably decent societies will tend to form anti-exploitation moral intuitions that conflict with act consequentialism. The claim is that if we reflect carefully we will come to be alienated in thought somewhat from these dispositions to judgment. On reflection, when we notice that ideals of fairness and reciprocity that have a pull on our loyalty lack generality, we question their status, and ultimately favor the ones that in context appear to be ones that function to bring about good consequences. Since rules serviceable for guiding choice and serving to coordinate people's actions will be coarse-grained, we will inevitably find cases where the exploitation norm applies but the circumstances that render it functional are absent. In these test cases we may well find our reflective judgments explained by act consequentialism.

The cases on which to focus are ones in which anti-exploitation norms and act consequentialism seem clearly to be opposed. Consider

CANNON FODDER. You find yourself in possession of political power. You conscript young individuals and coerce them into joining a military force that is waging a just war. The war will bring about morally valuable consequences, so the act of conscription is justified in act consequentialist terms, but the conscripted individuals do not see it this way. They harbor grave doubts about the justice of the war enterprise. You are inducing them to fight and die for goals that they reject and you uphold.

Consider also

HARD BARGAINING BY THE POOR. You are very badly off, with very grim life prospects. By chance you find yourself in a position of monopoly power. You are uniquely placed to rescue some very well off individuals, who have already sustained very high quality lives, from some peril. The peril might be slight, as when you are selling sunglasses to some wealthy tourists who are prone to succumbing to headache in the day's bright sunlight, or extreme, as when the peril is swift death. In either case, you could simply act to rescue the imperiled individuals at some reasonable (not excessive) cost and risk to yourself without demanding payment. You could demand some modest reward for rescue. But what you actually do is exploit your bargaining advantage as far as you can, gouging those whom you rescue by requiring them to pay a huge amount that is disproportionate to the cost you are incurring, though still one they are willing, in their desperate straits, to pay. By the correct act consequentialist moral principle, what they give up as the price of rescue generates more moral value in your hands, so if your only options are rescuing on fair terms or rescuing preceded by hard bargaining leading to what looks to be exorbitant price for rescue, act consequentialism endorses the latter course of action and condemns the former.

Notice that we would need to work hard to interpret the examples in such a way that it becomes plausible that the exploitative action really would produce best consequences and so be endorsed by act consequentialism as morally right. Coercing individuals to risk their lives for a cause they do not accept is not likely to lead to a spirited effort to attain the set goal and bring about good consequences. So we have to be imagining a special case with special features. Much the same is true of HARD BARGAINING BY THE POOR. In many settings, trying to gouge especially high profit for oneself in interaction with individuals who are socially better placed and more powerful than one is will likely lead to breakdown of cooperation with individuals from whose continued cooperation one could gain a lot. Moreover, since gouging the rich in this instance may well lead other rich people to decline to deal with other poor people when the rich can walk away from interaction with impunity, one may well be letting down the side, imposing costs on badly off people in the future by one's grabby action now. Socially powerful people may manage to find niches of power that enable them to gouge their interaction partners over the long haul, but here we are considering endorsing exploitation of the rich by the poor—not likely to be good for one in the medium and long term.

I submit that the expectation that hard bargaining that looks selfish and exploitative will not work out to one's advantage in cases like HARD BARGAINING and that manipulation of

people with an eye to inducing their sustained self-sacrifice in cases like CANNON FODDER is likely to go awry casts a shadow on our moral judgment of the cases. Being told that in these special cases all goes well does not fully erase the shadow. Moreover, we do tend to have a generic intuition against exploitation, and at least part of the explanation of our having this intuition is that the normal functioning of the norm against exploitation tends to be desirable by act consequentialist assessment. Taking these considerations on board in considering our response, we arguably will see the normative reach of the norm against exploitation as derivative and as having no weight in conflict with act consequentialist assessment.

The counterargument will be that we ought not to exploit just when exploiting gains us a bit better outcome. Exploitation is then a wrong-making feature of actions, a moral reason against a proposed course of action that has weight against the calculation of consequences. But perhaps the force of this claim should be blunted by the familiar observation that the advocate of multi-level consequentialism has an explanation for the moral appearances. A norm against exploitation that we ought to establish and sustain according to act consequentialism will not in its own terms allow exploitation for marginal gains because, humans being what they are, the consequences of establishing and sustaining such a rule would be less than optimal. This is to affirm act consequentialism not any version of rule consequentialism. But much of what Brad Hooker says by way of pointing out the consequentialist merits of accepting and internalizing moral rules of the W.D. Ross variety is plausible and can be affirmed by the act consequentialist (Hooker 2000).

Act consequentialism then can explain the agreed intuition stated near the start of this essay. But after reflection we may still find that it is counterintuitive to hold that there is really nothing per se morally wrong, at the fundamental level of principle, with taking unfair advantage of others when one is interacting with them. This residue of counterintuitiveness is a point against act consequentialism, but it is to be expected that every moral theory will be counterintuitive in some respects; this is compatible with deciding at the end of the day that the theory beats all competitors and is likely correct. (Of course this essay proposes no specific candidate act consequentialist principle. I am just issuing a promissory note.)

8. A puzzle solved.

A puzzling feature of common-sense moral thinking regarding exploitation is that exploiting people can turn out to be morally worse than having nothing to do with them even though all affected parties gain from the exploitation and both exploiter and victim voluntarily consent to the transaction. Consider a case in which the personal prerogative to pursue one's own projects and aims when one could do more good by helping others brings it about that it is morally permissible to do nothing for Smith on some occasion when he is in trouble and one could help and the moral norm against exploitation rules out as morally impermissible interacting with Smith by mutual agreement in a way that insists on exorbitant benefits for oneself and small benefit for Smith. Yet both parties to the transaction prefer interacting on these lopsided terms to having no dealings with each other. Common-sense morality is compatible with the further claim that it would be morally better to interact with Smith on more generous rather than these less generous terms. But it is still hard to swallow the idea that it is morally worse to help by unfair trade than not to help at all.

The multi-level consequentialist treatment of exploitation dissolves this puzzle. If the norm against exploitation that is accepted in society is the one that would be ideal from the consequentialist standpoint, then establishing and sustaining the norm produce better consequences than feasible alternative acts including acts directed to gaining acceptance for more fine-grained sophisticated rules. (For a different response to this same puzzle, see Pogge 2008.) Bringing it about that this rule is accepted and thus internalized by members of society predictably leads them to follow the rule on some occasions when doing so does not produce the best feasible outcome and is morally wrong. There is no inconsistency here. At the levels of

public morality and social norms, exploiting Smith by lending him money he needs at an exorbitant rate of interest is condemned. At the level of fundamental morality, declining to deal with Smith in the specified circumstances is more wrong than dealing with him on exploitative terms (though less wrong than dealing with him on more generous terms). The same might hold for the moral norm that accords everyone a personal prerogative to pursue her own projects, up to a point, even though this leads to worse outcomes. In a sense, the act consequentialist then supports a common-sense morality that countenances Pareto inefficiency. But the support is justified because and insofar as it serves long-run moral goals of bringing about morally better outcomes, and the support does not extend so far as to endorse the puzzling and objectionable claim that it is morally less wrong for you to exploit Smith than to have nothing at all to do with him in the imagined circumstances.

But why should we suppose that establishing and sustaining norms against excessive profit-taking when one happens to have bargaining leverage will tend overall to produce good outcomes even though in some circumstances they can discourage mutually beneficial trades? The surmise here is that establishing and sustaining such norms will tend to foster more deals in which those with bargaining leverage refrain from using it to extract a greater share of the cooperative surplus and that the overall consequences of such deals are sufficiently better in consequentialist terms that establishing and sustaining the norms are right choices from the act consequentialist perspective. This might be so even though the norms will predictably induce some to choose some acts that are really, not just in the short-term but in the long-term and all things considered, generative of worse consequences than might have been obtained and hence are morally wrong by the act consequentialist principle. (If this surmise should prove false, or false in some settings, the rapprochement between norms against exploitation and act consequentialism comes to an end. I bet the surmise is correct.)

References.

- Cohen, G., 2000, *If You're An Egalitarian, How Come You're So Rich?*, Harvard University Press, Cambridge, MA.
- Dancy, J., 2004, *Ethics without Principles*, Oxford University Press, Oxford.
- Dworkin, R., 2000, *Sovereign Virtue: The Theory and Practice of Equality*, Harvard University Press, Cambridge, MA.
- Hare, R.M., 1981, *Moral Thinking: Its Method, Levels, and Point*, Oxford University Press, Oxford.
- Hooker, Brad, 2000, *Ideal Code, Real World: A Rule Consequentialist Theory of Morality*, Oxford University Press, Oxford.
- Mill, J. S., 1978, *On Liberty*, ed. E. Rapaport, Hackett Publishing, Indianapolis, IN. Originally published 1859.
- Mill, J. S., 2001, *Utilitarianism*, ed. G. Sher, Hackett Publishing, Indianapolis, IN. Originally published 1861.
- Parfit, D., 1984, 'What Makes Someone's Life Go Best?', in his *Reasons and Persons*, Oxford University Press, Oxford, pp. 493-502.
- Parfit, D., 2000, 'Equality or Priority?', reprinted in Clayton, A., and Williams, A., eds. *The Ideal of Equality*, Macmillan and St. Martin's Press, New York, pp. 81-125.
- Pogge, T., 2008, 'Testing Our Drugs on the Poor Abroad,' in Emmanuel, E., and Hawkins, J., eds., *Exploitation and Developing Countries: The Ethics of Clinical Research*, Princeton University Press, Oxford and Princeton, pp. 105-141.
- Quinn, W., 1993, 'Actions, Intentions, and Consequences: The Doctrine of Double Effect,' reprinted in his collection *Morality and Action*, Cambridge University Press, Cambridge, pp. 175-193.
- Railton, P., 1984, 'Alienation, Consequentialism, and the Demands of Morality,' *Philosophy and Public Affairs*, vol. 13, no. 2, pp. 134-171.

- Ross, W. D., 1988, *The Right and the Good*, Hackett Publishing Co., Indianapolis, IN. First published 1930.
- Schmidtz, 2006, *The Elements of Justice*, Cambridge University Press, Cambridge.
- Sidgwick, H., 1907, *The Methods of Ethics*, 7th edn., Macmillan, London. First published 1874.
- Van Donselaar, Giis, 2009, *The Right to Exploit*, Oxford University Press, Oxford.
- Van Parijs, Philippe, 1995, *Real Freedom for All: What (if Anything) Can Justify Capitalism?*, Oxford University Press, Oxford.
- Wertheimer, Alan, 1996, *Exploitation*, Princeton University Press, Princeton.
- White, S., 2003, *The Civic Minimum: On the Rights and Obligations of Economic Citizenship*, Oxford University Press, Oxford.
- Wood, A., 1995, 'Exploitation,' *Social Philosophy and Policy*, vol. 12, no. 2, pp. 136-158.

About the Author.

Richard Arneson is a professor in the Department of Philosophy at UCSD. He holds the Valtz Family Chair in Philosophy. He works and writes on moral and political philosophy, with some focus on theories of justice.

-
- ¹ . Of course, the idea of act consequentialism is a shell that can be filled in with various contents. In the end the intuitive plausibility of act consequentialism stands or falls with the intuitive plausibility of its best version, with the best substance filling the shell.
- ² . The reader might reasonably protest at the claim in the text. I deny there are norms of fairness we should embrace and then I allow myself to argue by appeal to what is “uncontroversially fair.” My position is that we do have intuitions about what is and is not deontologically fair and your preferred deontological theory had better conform to them. As for myself, being a defender of act consequentialism, I deny any notion of deontological fairness has integrity and plausibility and should command conviction. A further qualification: I think there are simple fairness ideas that enter into the best account of the outcome assessment standard that specifies what is morally valuable. In a slogan: bringing about best consequences is bringing about good lives for people, with good fairly distributed. We should maximize a function of aggregate good that assigns special value to achieving gains for the deserving, and gains for the badly off. The special value claims are fairness claims. To get a determinate principle, these fairness ideas need to be specified in descriptive terms, with weights specified for the claimed priority relations.