

A NEW ARGUMENT FOR THE MAXIMIN PRINCIPLE¹

The natural distribution [of talents] is neither just nor unjust; nor is it unjust that persons are born into society at some particular position. These are simply natural facts. What is just and unjust is the way institutions deal with these facts. . . [T]here is no necessity for men to resign themselves to these contingencies. The social system is not an unchangeable order beyond human control but a pattern of human action. - John Rawls, *Theory of Justice*, p. 102.

In this paper, I will explore a new strategy of argument for John Rawls's conclusion that justice requires that primary social goods be distributed in accordance with a maximin principle, requiring that the least representative share be as great as possible.² The strategy of argument is in keeping with the spirit of the paragraph quoted above, and Rawls has made many of the claims I use to support the strategy. Yet the argument is different from the one that Rawls explicitly endorses.

The main idea is that Rawls's general conception of justice follows from taking seriously a stringent though overrideable prima facie duty not to harm others. I first schematically show how the idea leads to a Rawlsian theory of justice by presenting a simplified deductive version of the argument. Each premise of the schematic argument is plausible or can be replaced with similar plausible premises supporting the same conclusion. That is in itself some reason to take the argument seriously, at least if you accept this sort of theory of *prima facie* duties. But I actually think that there is better reason to take the argument seriously – that it has resources to meet many standard objections to a Rawlsian position. When striving for reflective equilibrium in this domain, it counts in favor of the Rawlsian conception of justice that it coheres with a plausible underlying moral theory, and that this theory gives it the resources to meet the standard objections to the

Rawlsian view in such a way that the whole package has a unified rationale. Thus the final part of the paper tries to make the case that this sort of congruence in fact obtains.

Outline of the Simplified Deductive Argument

The following 5 step argument presents the main ideas in a somewhat simplified fashion.

(1) We have a stringent *prima facie* duty not to harm others.³

(2) When choosing between options which each harm someone, our stringent duty not to harm implies that we should minimize the harm we do to any particular individual.⁴

(3) Choosing modes of social organization is choosing between options each of which harms someone.

(4) Maximizing the minimum share of primary social goods (PSGs) comes closest to minimizing the individual harms done in choosing modes of social organization.

(5) Therefore, we should distribute PSGs so as to maximize the minimum share.

In the five sections of the paper that immediately follow I will argue that the premises are plausible and that where there are equally plausible competitors to some of them, these competitors could be used to construct a parallel argument with a similar upshot.

PREMISE 1: WE HAVE A STRINGENT *PRIMA FACIE* DUTY NOT TO HARM OTHERS.

For purposes of the argument here I'm going to presuppose the sort of moral framework advocated by W. D Ross and endorsed by a number of others since then.⁵ Ross postulated a plurality of grounds for moral obligations, grounds which were such that more than one could be applicable in any given choice situation. Any given situation might be such that one could have multiple and possibly conflicting obligations to act due to the presence of multiple such *prima facie* grounds for

obligation. In such cases, certain duties might override others, and since no duty always wins out in such conflicts, each one can be overridden by others.

Ross gives a relatively short list of the different sorts of prima facie obligations we might have. What matters for my purposes is not the length of the list, but that each duty can be overridden in certain circumstances by others, and that a relatively stringent duty not to harm is one of those duties. According to the kind of stringency I have in mind, it takes really compelling reasons to override this duty, and an overall greater balance of benefits than costs is not in and of itself sufficient to do so. Thus the sort of view I presuppose endorses the commonsense idea that normally one should not harm one person even to prevent worse things from happening to others. For example, common sense forbids killing one person to save the lives of several others, even where life expectancies are equal.⁶

The prima facie duty not to harm is a side constraint on permissible ways of acting. Consequentialists famously deny that there are such constraints. They can make their rejection attractive by raising the costs of abiding by such constraints; many people become skeptical of this sort of constraint if it prevents us from avoiding huge and horrible costs. The attraction of conceiving of this constraint as merely a prima facie duty rather than an absolute duty is precisely that it allows us to defend against such “cost-raising” arguments. A prima facie principle requires only that it takes significantly more good to justify harming a person than it takes to justify allowing harm to come to that person.⁷ Thus, the objection that constraints make it impossible to avoid even the greatest of disasters does not count against this constraint, even if such objections would have some force against absolutist side constraints.

Certain nonconsequentialists will dislike the idea of a duty not to harm when harm is

construed in the way I am thinking of it. They will claim that there is a duty not to harm people only in a moralized sense, where harm is construed as making a person off by taking away something the person is entitled to or has a right to. Such non-consequentialists think that my argument for the maximin principle equivocates on the notion of harm. The first premise is true, they think, only if harm means doing something to a person that adversely deprives them of something they have a right to. But, the remaining premises require a notion of harm on which harm occurs wherever someone makes someone else worse off by what they do, whether one is depriving a person of a right or not.⁸

I need not and do not deny that people have a duty not to deprive others of what they are justly entitled to.⁹ I claim only that there is also a duty not to harm others that does not depend on a prior notion of what rights people have. And the argument here employs that notion.

There are of course other serious objections to the Ross-inspired framework which includes defeasible constraints, insofar as there are competing plausible moral theories. But for my purposes it is enough to rely on these brief rebuttals and on the work of others for more detailed defenses of this sort of theory. If the argument that results from using the theory as I do provides support for the Rawlsian conclusions that I claim, it will reinforce the existing arguments by showing that they can be mutually supporting components of a comprehensive view which survives the process of reaching reflective equilibrium about morality and justice.

PREMISE 2: WHEN CHOOSING BETWEEN OPTIONS WHICH EACH HARM SOMEONE, OUR STRINGENT DUTY NOT TO HARM IMPLIES THAT WE SHOULD MINIMIZE THE HARM WE DO TO ANY PARTICULAR INDIVIDUAL.

A distinctive feature of prima facie duties is that they will conflict with other prima facie

duties and that no prima facie duty is such as to always win out in such conflicts. Even if there were no other prima facie duties besides the duty not to harm, it looks like a person could wind up in a situation where the duty conflicts with itself. Perhaps it is possible that someone who never did anything at all could harm no one. But once we live our lives for a while we will put ourselves in situations where each of our options involves harming someone or other. Examples abound in the vast literature on the "trolley problem,"¹⁰ and they arise in real life as well. It is thus worth thinking about whether there might be principles for deciding what to do in conflict cases, at least for certain domains.¹¹ And since the kinds of conflict that are most relevant to my argument here all involve the duty not to harm conflicting with itself, I'm going to stick to canvassing principles for situations in which every course of action harms someone or other.

One candidate principle for such situations prescribes minimizing the harm we do to any given person. Call this the "Minimax Harm Principle". This is the principle I use in the simplified argument for Rawlsian principles that I give above. I think it fits well with the idea of a side-constraint against harming others, but there are other possible principles for handling conflicts of harm which also seem to be in that spirit. One such principle might require harming those who are already very badly off, as little as possible. Intuitively it seems to me to have many of the same attractions as the minimax harm principle, and on certain conceptions of an individual's well-being they might amount to the same thing.¹² I think that a revised version of the main argument would survive substituting such principles for the Minimax Harm Principle, so I will not argue against this competitor to Minimax Harm.

Other competitors would however not allow the overall argument to go through. The most obvious such competitor would require us to minimize overall harms. That is what consequentialism

would require in domains where the choice between options is solely an issue of deciding between competing actions all of which harm someone, and where no side benefits are involved. But since we have already abandoned consequentialism in admitting that one has a more stringent duty to avoid harming anyone, minimizing overall harm has no special presumption in its favor. Minimizing overall harm and Minimax Harm will agree on cases where at most one person is harmed. Disagreement only shows up in cases where harms will effect larger numbers, for example where we can avoid harming a smaller number of people quite a lot only by harming some greater number of people somewhat less. In such situations, the following considerations seem to me to count in favor of Minimax Harm as opposed to minimizing overall harm.

If you find duties not to harm attractive, it seems plausible that they should be more stringent the greater the harm suffered by the person harmed. We can see this when we look at what it would take to defeat the duty. Most people think that the duty not to harm may sometimes be over-ridden in order to defend oneself or others. But the appropriate measures depend in part on what we are defending against. It may be permissible to slap or kick a person engaged in a non-lethal attack on an innocent person, but not permissible to kill them. And I find it implausible that there is some number of people subjected to non-lethal attack which would make killing permissible after all.¹³ On the other hand, where a lethal attack on an innocent person is involved, it is permissible to kill more than one attacker to save the victim. What it takes to defeat the duty not to harm for small harms and what it takes for much greater harms seems to be a distinction of kind.

If the foregoing is correct, the cases in which Minimax Harm gives guidance all involve conflicts between more and less stringent duties. Are there principles for deciding when more stringent duties conflict with greater numbers of less stringent duties? Nonconsequentialist fans of

the duty not to harm have already committed to ruling one way of deciding for one kind of conflict. Where meeting one duty requires harming another, and the other requires only not-helping, it is not just a matter of weighing up the total benefits on either side. While there may be a threshold beyond which this no longer holds, up to that threshold the duty not to harm wins out.

The minimax harm principle applied to conflicting actions each of which harms someone treats all conflicts between more and less stringent duties in a similar way. In principle a sufficient sum of lesser harms adding up to much greater overall harm could justify us in violating the Minimax harm Principle. When that happens we now have a decision to make about who to harm. The argument above suggests that this conflict can be cast as one between a more stringent duty not to do greater harms to any given person, and a less stringent duty not to cause lesser harms. It is more continuous with the treatment of such conflicts in the harming/not-helping case to allow the more stringent duty to win out past the point where the balance of total harm is greater because the greater stringency of the duty not to cause great individual harm requires more than just a lower aggregate of badness avoided to justify overriding that stringent requirement.

That continuity will make many justifications for a stringent duty not to harm more compatible with the Minimax Harm Principle than with the alternative weighing up of total harms. For example, one line of defense of the duty not to harm is to stress that people differ from other objects in having a subjectivity, that is a point of view on what happens to them. Respect for people requires taking that point of view into account if we are going to allow our plans to affect them in significant ways.¹⁴ Where a person's interest loses out in a balancing test, we ignore it, except in a certain counterfactual way. We ignore it insofar as it does not modify our behavior in pursuing our ends, though it is true that had the balance of goods on either side been more equal, it could have

affected our behavior by tipping the balance.

The same rationale will have implications when choosing between actions each of which harm someone. If we treat the decision as one of balancing total harms, we are ignoring the interests of those who lose out in the same way as we do when we harm one to help more. Can we do otherwise by deciding otherwise? No and yes. Even acting as Minimax Harm would require we will not allow us to refrain from harming each and every person we effect. But we will never be permitted to ignore a person's interest come what may, just because there are so many others getting relief at their expense. For each person, once their interest is impacted to a degree greater than anyone else's, they will have priority.

There is much more to say here. Similar considerations could be offered on behalf of a principle which required us harm those already badly off as little as possible. But as I already have claimed, such a principle will function similarly to the minimax harm principle in the overall argument. Both of these distribution -ensitive principles fit better with the vaguely Kantian respect-based justification of side constraints such as these invoked immediately above than their non-distribution sensitive competitors do.

On a less abstract level, there are other, more particular judgements, that support the Minimax Harm Principle as well. For example, it may help to think about how we would go about compensating people for past harms done to them. Common sense says that other things being equal, we should compensate those who have been harmed most before those harmed least. Where things are not equal because it will cost more to compensate those harmed the most so that we can do less total compensating if we concentrate on those harmed most, it still seems that they should get some priority unless the benefits due to efficiency are very very great. A society has the wrong

priorities if it ignores the effects of systemic racial injustice while ensuring that property owners are compensated for economic losses due to zoning decisions. Thus it seems intuitively plausible that severe harms should have priority where compensation is concerned.

Compensation for harms functions to make things as nearly right as we can make it. It mitigates the harmful effects of harmful actions, and where compensation is complete, eliminates those effects. Full compensation for a harmful action makes for an overall course of action from which no harm results. When we cannot fully compensate everyone, we should give priority to those harmed most. This indicates that an overall course of action in which the maximal harms are as small as possible is to be preferred over alternatives. And if that priority is to remain, even at some cost to efficiency so that we can do less compensating overall, we have reason to believe that minimizing overall harm is less important than minimizing harm to any given person. Considerations such as this favor distribution sensitive principles for allocating harm over those that worry only about aggregate amounts of harm. The minimax harm principle is distribution sensitive but we can keep other distribution-sensitive competitors in mind as we fill out the rest of the argument.

**PREMISE 3: CHOOSING MODES OF SOCIAL ORGANIZATION IS CHOOSING
BETWEEN OPTIONS EACH OF WHICH HARMS SOMEONE.¹⁵**

I have already said that the duty not to harm can conflict with itself. Politics, or political economy, is full of such dilemmas. We build irrigation projects and flood control devices, regulating the flow of water for the presumed benefit of all, only to have these projects and policies exacerbate the disastrous effects of heavy weather, leaving us with choices about whose homes to flood. And, these are not unusual cases. The problem of distributing the benefits of social

cooperation is itself a problem of this kind, as I'll argue below. As a result, we will wish to have something to say about distributive issues of this sort.

Shortly, I will argue that a maximin principle like that advocated by Rawls will serve the needed purpose, at least for the fundamental political cases, but first I wish to expand on the point that the problem of distributing social goods, is a problem of conflicting obligations not to harm others. The basic thought here is that any distribution of resources to some people will have opportunity costs for other people. While the distribution of resources is not a zero sum game, every distribution leaves some worse off than they might be under an available alternative. Depending on the baseline from which opportunity costs should be measured for the purposes of determining harm, putting such costs on another may well count as harming that person.

The welfare of each of us obviously depends on a vast network of social cooperation, calling upon the actions and the resources of a great number of other people. To begin with, each of us was born an infant in no position to care for ourselves. But also few of us could by ourselves produce the necessary food, shelter, and so on to keep ourselves alive without relying on the cooperation of others at some stage. It is no exaggeration to say that without the help of other human beings, we would all be dead. Thus, relative to a state of no cooperation, we all benefit from cooperative social arrangements.

On the other hand, for almost all of us, other imaginable and possible social arrangements would leave us materially better off than we are now. Imagine a world in which you have the wealth of Gates or Buffet.¹⁶ Or imagine a world with more egalitarian distributions than ours, where most people do better than most people currently do. These worlds, and a good many more, would be possible if different institutions and conventions governed human interaction. And, institutions and

conventions require our cooperation to function as they do - not in the sense that any one of us is singly able to change them, but in the sense that we could collectively do so.

This means that whether each of us does well, or has an opportunity to do well, depends on which social arrangement comes into being through the confluence of all the various actions people take and have taken. How well I do depends on whether the society I live in favors those with my talents and/or disfavors those with my disabilities. In fact, speaking strictly, it is often not possible to label a trait as a talent or disability until we know the society to which the person with the trait belongs. Thus, our actions which support these social arrangements have the potential either to help or harm various individuals, depending on which arrangements they support, and which traits those arrangements favor.

But, it is obvious that there is no one world, in which each of us is maximally well off, or has the capability to be maximally well off. Maximally expanding the opportunities for one person, by necessity limits the opportunities of some others. Thus, benefits for some can only come at the expense of others. This does not yet show that every society will benefit some at the cost of harming others. For, being harmed involves more than just not doing as well as one could. What counts as a harm and what counts merely as a refusal to help is relative to some sort of baseline. But, I will argue that **some** instances where people are prevented from doing as well as they can do count as harms.

**Depriving a Person of Opportunities That They Would Have Had Under An Alternative
Arrangement Is Sometimes Harming That Person**

One can harm a person by depriving that person of an opportunity of which they have not yet availed themselves. Many of the things that are most important to us, are important not because

of their current effects, but because they enhance future possibilities. As a corollary, depriving a person of such chances counts as a harm. A chance for a decent education can be a benefit to someone, even when that person does not enjoy schooling, because that education might enhance their prospects in other ways. If a mail robbery causes my application for admission to be lost, I have lost my chance for admission and been harmed. Perhaps no harm was done if I had no chance of being admitted. But it is not enough merely to show that I was not admitted yet when the envelope was purloined. Nor would it be enough that what I regard as the benefits of a decent education (good employment, etc.) would not accrue to me until long after I graduated.

Some of the greatest harms that befall us do not change the status quo, in the sense of putting us in a new situation, except insofar as they make it less likely for us to change the status quo. Many of these are political in nature. When Martin Luther King was assassinated, it harmed not just those close to him, but also those whose aspirations were supported by his leadership in the civil rights movement. When political reforms fail to get necessary political support in congress, their failure is often said to harm those whom they would benefit. Depriving a person of the benefits that she would have received under an alternative sociopolitical arrangement can in principle count as harming her in just this way.

What is or is not a harm will depend on what we take to be the legitimate baseline for comparison. But even without agreeing on a baseline, I think that we can agree that certain political and economic arrangements harm those who are worst off under those arrangements, even when they do not reduce those people to the level of a "state of nature," that is to the level they would have been at in some "pre-social" non-cooperative state. For example, it seems obvious that slavery involves slave-holders and any who act to support the institution in harming slaves. And we should

not change our judgement if we discover that slaves are offered the option of permanent exile away from civilization. For the baseline relative to which an arrangement counts as harming a person must be sensitive to the other arrangements possible but not actual. And among those arrangements would be a society in which there were no slaves, and in which the benefits and burdens of cooperation were more equitably distributed.¹⁷ This shows that depriving a person of alternative modes of social cooperation can constitute harm, even if those persons are left better off than they would be without social cooperation.

The reason why each possible arrangement counts as an institution which harms others, as opposed to merely refusing to help, is in part that each such arrangement significantly limits our prospects. Each institution has opportunity costs for some people, and for those bearing the brunt of those costs the institution can well constitute a harm to their interests. When we codify laws establishing property rights in objects, we limit the opportunities for using those objects that might have been had by others.¹⁸ Those limits may or may not count as harms depending on how much worse off than a reasonable baseline those limits leave those limited. But those limits are clearly a result of our agency and hence count as doings of ours.

Which baseline?

We are now faced with figuring out how such alternative arrangements determine what an appropriate baseline for determining harms might be. It is easy enough to say that the baseline rules a harm any outcome which leaves a person as badly off as they would have been without social cooperation; the examples above show that. It is harder to draw the line among outcomes all of which leave no one worse off than they would be with no social cooperation. We cannot say that outcomes which leave a person with less than they are entitled to, harm the person. For, though this

may be correct, we are trying to figure out what a person is entitled to by trying to figure out what counts as harming them.

The outcomes towards the middle of the range of possibilities seem most promising. Once we realize that depriving someone of a possible opportunity can harm that person, we are committed to making the baseline for determining harm sensitive to non-actual possibilities. Restricting a person's options so that what they are left with are those possibilities that are the least attractive of those antecedently possible harms them. But eliminating some of the very best while leaving the person in a situation that is as good as the average likely option prior to the restriction of options does not seem to. This suggests that the relevant baseline is a one which averages a person's prospects across all possible alternatives weighted by their likelihood.

Another way to approach the determination of an appropriate baseline given multiple possibilities reinforces the general idea. To harm someone is to worsen their position through the sort of agency that an appropriate account of the doing allowing distinction would highlight. Thus, we should think of what outcomes might worsen a person's position in a situation where a number of outcomes are antecedently possible. If narrowing the range of possibilities to some particular outcome counts as a bad thing from that person's point of view, we have reason to think that the outcome is below the relevant baseline. Suppose you buy a ticket for a lottery where there are ten possible distinct outcomes each paying out a multiple of one hundred dollars, covering the range from one hundred to one thousand dollars, depending on the outcome of a sequence of truly random events. Each outcome is equally likely, so that the expected benefit of buying the ticket is five hundred and fifty dollars. If the outcome of the lottery is seven hundred dollars, this will be good news from your point of view and you will be better off after the lottery has been determined than

you were previously. Similarly, if the outcome is four hundred dollars, that will be bad news from your point of view and you will reasonably feel that you are worse off than before that outcome was determined.¹⁹ So the baseline for determining when a narrowing of the range of possible outcomes makes you worse off is a function of the values of the various alternatives, and of their likelihood.²⁰

One such baseline might be determined by something like the expected utility for any person averaged across all the realistically possible outcomes, taking their probabilities into account. An argument for such a baseline might go as follows: If each of the possible alternative arrangements were solely the product of choice, we could think of each alternative as equally likely. A person's current well-being could be seen as a function of that person's expected well-being in each of the practicable alternative arrangements multiplied by its probability.²¹ We could then use the average well-being for that person across the possible social arrangements as our baseline for determining harm, or use a package of resources sufficient to give the person that amount of well-being. If we disallow arrangements that make some worse off for no benefit to anyone else, the resulting average will be neither at the top nor the bottom of the range.²² In effect we are thinking of each of the alternatives as having opportunity costs for every agent. When the opportunity costs are large enough to put the remaining alternatives for an agent below the average for all the antecedently possible or likely options, we will say the alternative harms that agent.²³

Any Set of Institutions Will Leave Some Below The Baseline

Given a reasonable baseline for determining the harms done in choosing a set of distributive institutions, any choice will likely harm someone. Compared to the sort of expectational baseline illustrated above, any likely actual social arrangement will harm some who find themselves in it.

For, it is unlikely that there will be even one in which all do as well as the average of their expectation across all likely possibilities. Most people's expected utility computed as a function of the possible social arrangements and factoring in the probabilities of each will be neither extremely high, nor extremely low. Yet even relatively egalitarian social arrangements leave some people rather badly off, significantly worse off than they might have been under many other realistically possible arrangements. When we average in all of these possibilities, the worst off in even such egalitarian societies are unlikely to do as well.

I suspect we would have a similar result using other plausible candidates for a baseline. Taking some more objective measure of well-being, say expectation of primary social goods, would not be likely to yield a baseline low enough that realistically possible social systems would leave no one below this baseline. Nor would playing around with methods other than averaging expectations however measured. The only plausible method I can think of that might avoid the conclusion is allowing the maximal minimum share (measured in either resources or well-being) to determine the baseline. For it would be possible to create a society in which the worst off member was as well off as possible. Such a baseline would have the virtue of allowing me to construct an even shorter argument from the duty not to harm to my conclusion. But it does have some drawbacks. It isn't obvious that the worst off in such a society are not harmed, in virtue of being as badly off as they are. And, as I have argued, a baseline for harming must be sensitive to the other possibilities open to a person. A society which maximized the minimum share of primary goods or whatever, might give that minimum share to someone who would in all likelihood be better off under most alternative arrangements. If so, it seems that she was harmed by the choice of this one. I conclude, that the relevant baseline for determining harms will have to be set somewhat higher than

the minimum set by a maximin principle. And given such higher baselines, we should conclude that each actual social arrangement we might choose would harm someone.

Even if I am right about that, someone might wish to object that this is not enough to show that we, meaning you and I, must harm someone. For, couldn't we by refusing to participate in politics, avoid acting in any way which causes the social arrangements which in turn cause harm to some? The line between harming and merely not helping is one which tracks our agency in bringing about an outcome; couldn't we avoid harming by not acting?

To this I have two replies. The first is just that we have already acted in ways that underwrite current social arrangements. Most of us, by the time we come to consider political arguments of the kind I am making, have done our part at various points in our lives to support existing institutions. We are more like the pilot of a plane which will crash and kill some, now having control only over how many she will kill, than like the pilot at any time before that plane has left the ground. Given what we have already done, we will harm someone.

Second, even if we weren't in this position, and even if we could abstain from social cooperation to avoid harming others, we could only do so with grave consequences for our own well-being. It wouldn't be enough to stop participating in politics. We would also have to stop most economic activity insofar as our economic choices change the legal rights and remedies available to others, and set up incentives that have global effects. In order to truly make social action not a result of our agency, we would have to withdraw radically from economic and political life. Perhaps we could do this and survive. But I'm not sure that we are morally required to do this, even by our stringent duty not to harm. Among our duties we seem to have a prerogative to take care of our own well-being provided it does not extract too high a price from others. I don't want to call this a duty

to oneself. But it would seem to me to overstate the stringency of our obligation not to harm others to say that it always over-rides any consideration of one's own interest, however pressing.

If that's right, a prerogative to look out for one's own important interests could justify one in violating our duty not to harm if need be. But we should still try to minimize our violation of that duty. And that will involve choosing between alternative courses of action that allow one to serve one's own important interests, but which also each harm someone or other. This brings us back to the problem discussed earlier, that politics and political economy seem to essentially involve choosing whom to harm. We need a principle for deciding between such competing harms.

**PREMISE 4: MAXIMIZING THE MINIMUM SHARE OF PRIMARY SOCIAL
GOODS COMES CLOSEST TO MINIMIZING THE INDIVIDUAL HARMS DONE
IN CHOOSING MODES OF SOCIAL ORGANIZATION.**

My argument in this section will have two parts. First, I will argue that nonconsequentialists should use expectation of primary social goods as the measure of harms to the expected well-being of others, where people have similar capacities. Then, I will argue that the minimax harm principle applied to social choice situations will yield the maximin principle using primary social goods as a measure.

Primary Social Goods

Within a non-consequentialist framework of the sort that makes much of the distinction between not helping and harming, the use of primary social goods as a measure of comparison makes a good deal of sense, at least for those without severe handicaps. Among the bad things that happen to others that will not count as my harming them, are those bad things they do to

themselves.²⁴ If I give someone some food, and they let it spoil rather than put it away, their hunger is not a doing of mine. When I leave a person with a bunch of resources which would be enough to do a number of things that they choose not to do, any harm that results from their choice is not something I have done to them. If I have harmed them, it will be because the resources I have left them with are not usable to the extent necessary to leave them as well off as the resources they could previously expect would allow them to be.

Distributing a package of primary social goods to a person typically gives them a range of available options. Wealth and income, liberties, powers and prerogatives of office, are things which allow people of relatively normal abilities to achieve a certain level of capability or well-being. They may or may not actually use them so as to benefit themselves, but so long as there is no real impediment to their so using these goods, any harm from their not being used is their own doing. Where people have handicaps that prevent them from using these resources to their benefit, that may not be the case. Such people may need a different package of goods to achieve similar well-being or capacity for well-being, and we might be harming them, if we leave them without the special resources they would need. But for the more usual cases, primary social goods will be all that is necessary to give people the potential for a significant range of capabilities.²⁵ So long as we allow them enough primary goods so that their prospects are roughly what they would be under the relevant baseline for computing harm, we have not harmed them.

Minimax Harm Yields Maximin PSGs

If this is correct and if there is a baseline for computing harm, a person will be harmed by a social arrangement if the package of primary social goods distributed to them by that arrangement is less attractive than a position at the baseline. The extent of the harm will be proportional to the

difference in attractiveness of the two situations. Thus the minimax harm principle will require us to minimize the maximum disadvantageous difference between a package of goods sufficient to reach the baseline and the actual package that any person will be forced to tolerate.

Of the distributive principles that can be relatively easily applied, maximin primary social goods is most likely to approximate this goal. That is because the different baseline expectations for different people are unlikely to differ by all that much. I have argued above that most people could expect a baseline package somewhere near the middle of the range. If we adopt the simplifying assumption that their expectations are the same, maximin primary social goods would give just the right result. Maximizing the smallest representative package will have the effect of minimizing the largest deviation from the package of primary social goods that our harm baseline would require.

In fact, the simplifying assumption, that we each have the same expectation of PSGs when all the possible social arrangements are factored in, is not so far from the truth. Each of us could be a monarch, and each of us could live at a subsistence level. To the extent that I have some potentially lucrative talents you don't have, you likely have others that would be equally well remunerated under some alternative arrangements. Furthermore, many of the talents I actually have, are ones I might not have had in some alternative social arrangements, and some I don't have I might have had in a different social environment. Also relevant is that which traits actually are talents is a function of social arrangements, so that even if some propensities will be useful over a wide range of arrangements, the extent of their remuneration will vary greatly. All of this is consistent with some differences in average expectations of primary goods over the range of different social arrangements, but these considerations make it likely that for most of us those differences will not be large relative to the average.

Given the relatively small range of expectational variations, and given the epistemological problems with determining what the variations are, using the simplified common baseline has much to recommend it. Where we are unsure of how the variation goes, and concerned to minimize the maximum harm, caution would dictate not creating the possibility of greater maximum harms by assuming that we could allow some to be worse off than anybody really need be, on the assumption that we have guessed correctly that their average expectational package of PSGs was lower than most. Even what seem to be people's natural abilities under current arrangements are often results of those very arrangements. While there is very likely a genetic component in the sorts of skills that are likely to have value across social arrangements (such as perhaps some sorts of intelligence), social factors such as one's parents economic background very likely play a larger role. For example, being poor correlates rather strongly with restricted opportunities to develop one's intelligence while young, exposure to lead paint and other intelligence-stunting poisons, and so on. Determining with any degree of confidence which people have greater natural talents is a tricky business at best.

Someone who wanted to ignore the difficulties in an attempt to allow deviations from maximin where inherent natural abilities might set higher and lower baselines for potential harms, would need to come up with a mechanism for doing so. Two sorts seem possible and both are fraught with grave difficulties. One sort would allow direct judgements of a particular person's natural abilities to figure in justifying deviations from maximin. Someone (some person, some committee, some governmental authority?) would be charged with making determinations about a person's natural ability and computing the threshold below which their situation could count as harming them.

Given the uncertainty of the issue our judgements here will likely diverge. Political strife would be inevitable, to say nothing of the tremendous waste of resources in trying to make such

determinations and then seeing to it that each individual person's actual situation comported with the verdict. Simplifying by ignoring a small deviation from the ideal will avoid the sorts of political dispute that could make everyone less well off than adopting principles which force us to decide such issues.

The only alternative that would still purport to take differences in natural talents into account would be some sort of invisible hand mechanism. The idea would be that we could somehow construct markets or institutions that would distribute resources in proportion with "natural" talents, though no one would be charged with seeing to it that resources were so distributed. The prospects here seem equally dismal. Given basic disagreements on what a "natural" talent really is, and over who is so talented, it would be wishful thinking to suppose that this option could be used both to publicly legitimate the institutions in question, and avoid political strife. Furthermore, I see no reason to think that such institutions could be constructed. Any system of rewards tends to create niche markets, not for people who are more talented but for people who are able to figure out the system and figure out how to work it to their benefit. At the same time, many who are talented do not reap similar benefits only because their talent does not extend to working systems of rewards to their own benefit. It is this last part of the objection that is crucial given my overall line of argument. Those not untalented people at the bottom who end up with a share smaller than the maximin principle requires are being harmed to satisfy more tenuous claims by others that they would be harmed by accepting the constraints maximin imposes.

The simplifying assumption that we are all roughly equally naturally talented, and hence have about the same expectation of good across various social arrangements, is thus superior to the alternative of trying to factor in only speculatively determinable differences in natural talents.

**CONCLUSION: WE SHOULD DISTRIBUTE PRIMARY SOCIAL GOODS
SO AS TO MAXIMIZE THE MINIMUM SHARE.**

My initial argument is now complete. I have provided persuasive though not conclusive reasons for accepting each of its premises. I will now show that this argumentative strategy enables the proponent of maximin to answer a number of at least initially plausible objections, and to adjust the overall theory to handle the more telling of the objections in ways that are not *ad hoc*.

**HOW THE ARGUMENT ASSISTS US IN HANDLING SOME STANDARD
OBJECTIONS TO THE MAXIMIN PRINCIPLE**

Since Rawls first defended his theory of justice, it has been subjected to vigorous attack. In the sections that follow, I will defend the theory using the resources made available by the justification I have just offered. While a couple of the objections will lead us to allow for exceptions to the general maximin principle in a limited range of cases, it will turn out that in those cases, the not-harming rationale supports making the exceptions. In the case of other objections, the resources of the not-harming rationale will give us ways of denying the force of the objections in a non-arbitrary way. The overall idea is that the plausibility of the sustainable objections stems from the plausibility of the proffered non-consequentialist rationale. Adjustments needed to ensure that distributive principles track the underlying justification's prohibition on harming do not undermine the principle's authority for the remaining cases, provided that the rationale supports the principle for the rest of its domain.

Objection 1: Maximin Asks Too Much.

Utilitarian critics argue that the maximin principle asks too much, insofar as it would require forgoing great benefits to some in order to avoid worsening the lowest social position. And even

non-utilitarians may fear that the maximin principle asks too much. For, if social distribution necessarily harms some, could there not be cases where we could avoid great harms to many others, if only we would be willing to slightly worsen the worst off position? For example wouldn't it be permissible to slightly lower the income for the poor if it would significantly raise the income of those in the middle? Maximin would rule this out.

Reply

Using primary social goods as the measure of what is distributed makes the envisaged situation very unlikely. Take income for an example. Even if some inequalities increase the overall amount to be distributed, it is unlikely that making the bottom position less attractive yet would somehow create enough additional wealth to significantly improve the position of those in the next position. In most cases where there are two options, one of which makes those in a more advantaged position very much better off and those at the bottom a bit worse off than the other, there will very likely also be a third possibility somewhere between the two in which the worst off do better yet.²⁶

The conclusion is reinforced when we recall the usual explanation for how inequalities can generate additional income, by offering incentives to work harder. Who is supposed to receive those incentives in the cases posited by the objection? Perhaps those at the bottom will work harder in a (vain) attempt to get out of their sorry position. Is it really likely that a small worsening of their position would entice them to be significantly more productive? Perhaps those on the next rung of the economic ladder will work harder for fear of dropping to the position of the worst off. But again, why would a slight worsening of the position of the worst off lead to so much greater productivity? There is no reason to think that the kind of case imagined could ever arise.

As a response to a hypothetical objection offered in response to a moral principle, this may

seem *ad hoc*. It is not normally an answer to a hypothetical counter-example to a moral principle that the postulated example is not likely to happen. But this is where our defense of the maximin principle comes to the rescue. Supposing great benefits could be achieved for the better off at small expense to those who are less well off, we need not insist that the maximin principle be followed. We can allow that for great benefits it can be violated (as we have for the duty not to harm and the minimax harm principle), while relying on the argument above to show that such a situation will never arise. Since the maximin principle is based on the duty not to harm and the minimax harm principle as I have argued, when the latter is defeasible the former should be as well. In principle, if enough overall harm is avoided, it will be OK if the maximal harm be somewhat greater than strictly the minimum possible. That thought is no more plausible or implausible than the thought that the duty not to harm might be less than absolute. Thus, if it is reasonable to concede to the objection that we should have the least well off make a small sacrifice, it will also be reasonable to think of the not-harming obligation in a similar way. This response is not *ad hoc*.

The concession is in fact entailed by any non-absolutist reading of the underlying duty not to harm. But it is a concession that is merely hypothetical, if the original line of response - that such situations are unlikely - is correct. Allowing the concession for such hypothetical cases, does not undermine our sticking with maximin wherever the goods to be gotten by giving it up are not such as to over-ride the duty not to harm. And, the original line of response is designed to show precisely that in the actual world such cases never arise. Hence, in the real world, we should stick with the maximin principle.

Objection 2: Maximin Unfairly Favors the Poor over the Rich.

A second objection is in a way related to the first. This is Nozick's objection that while

Rawls may have shown that the poorest people in a society governed by the maximin principle have no grounds to complain since they are not asked to sacrifice for the sake of the better off, Rawls has not shown that the richer people have no such grounds. For, (the objection goes) the better off are being asked to sacrifice for the sake of the less well off.²⁷ The theory treats the rich and the poor asymmetrically, and hence unfairly. A utilitarian might endorse the same sort of criticism, arguing that maximin violates the principle that each person's welfare should count equally.

Reply

The proposed line of argument for the maximin principle gives us an immediate reply. It is just false that the better off are being asked to sacrifice for the sake of the less well off in a system governed by the maximin principle, though they are being asked to forgo additional benefits that the less well off might have to sacrifice to allow them. At least that is how things will look from a non-consequentialist perspective that distinguishes harming another from failure to help another. Equal treatment requires first that we harm no one, and only secondarily that we help those we can help to a similar degree. This can be denied by the utilitarian or other consequentialist, but to show that Rawls is making a mistake in not thinking of the better off as making an unfair sacrifice, the critic will need to show that the underlying non-consequentialist distinction is unsound. And that should be no comfort to the Nozickean, whose defense of rights as side-constraints presupposes just the sort of distinction we are employing.

Objection 3: If Maximin Is Correct, It Should Apply Universally.

There are two ways to take this objection. One way is to use it as a premise in a *reductio*, by finding a situation where it seems implausible to distribute resources as maximin would require. The other is to criticize Rawls's limitation of the scope of the principle to persons cooperating for

mutual benefit. The former position is adopted by Robert Nozick, when he ridicules the idea that more and less fortunate Robinson Crusoes have any redistributive duties to each other, while the second seems to be the position of Brian Barry.²⁸

Reply

Nozick may be right; a bunch of Robinson Crusoes marooned on different islands may have no duties to redistribute their resources for egalitarian purposes so long as they are genuinely stuck on separate islands. That is because they aren't harming each other by withholding such aid. But he is wrong that once they are capable of interacting (by getting off their islands and moving around) their obligations are not significantly changed. Given that a mobile but less fortunate Crusoe could now go to a better stocked island, any rival Crusoe who got there first and prevented him from taking advantage of those stocks might now be harming him. If so, our not-harming rationale may apply while it did not before.

Barry is wrong to at least this extent: It is not completely arbitrary of Rawls to limit maximin to cooperators rather than isolated Crusoes, insofar as a not harming rationale will apply to the former, but not the latter. There may of course be other reasons to redistribute even in the absence of cooperation, but that is another issue.

Objection 4: Maximin is Insufficiently Egalitarian

The above criticisms (excepting Barry's) came from the right, from those who fear that maximin gives too much weight to the interests of the less well off. Maximin has also been criticized from the left, by those who think it is not egalitarian enough. Maximin puts more emphasis on improving the position of the worst off than it does on achieving overall equality. In so doing, it downplays the importance of communal values such as fraternity, allowing some to improve their

own positions without doing all they can to improve the position of those less well off. When the size of the pool of social goods is fixed, an equal distribution is also the maximin distribution. The two principles come apart only if the amount of goods available for distribution can be increased by allowing inequalities. And this is most likely to occur only where incentives are necessary to get people to work harder or more efficiently. But if incentives are necessary for this end, that shows that some people are holding out (whether consciously or unconsciously) for a larger than equal share of the social product before they are willing to contribute to the greater support of their fellow citizens. Maximin thus panders to the baser self-interested motives of citizens, rather than to attitudes more compatible with an overall egalitarian theory.²⁹

Reply

Several points are worth making before giving a full reply. The first is just to note that the maximin principle only deviates from strict equality **if** such incentives are in fact necessary to gain enough cooperation so as to raise everyone's position above a baseline defined by equality. I myself think it unlikely that anything like the incentives now in place in Europe, let alone the United States, would be necessary to gain willing cooperation. It may even be that when things are more equally divided there are more incentives for increased cooperative production than under current circumstances.

Second, an advocate of the maximin principle should not be taken to deny that the cultivation of civic virtue in citizens is unrelated to issues of distributive justice. For the attitudes of citizens towards what justice requires of them will be important both to insuring overall stability, and also to the smooth functioning of principles so that they benefit the position of the least well off. Thus, if public application of maximin principles did so undermine the spirit of cooperation that it would not

be long before citizens failed to cooperate the principle itself would dictate that we move to a principle such as strict equality in order to most closely satisfy the maximin criterion over the long term. As the objection itself notes, the maximin principle deviates from strict equality only where an equal distribution leaves no one worse off than anyone under the equal distribution, and a good number do better yet.

These points noted, we can now draw on the non-consequentialist justification for the maximin principle to answer the strict egalitarian by noting that the duty not to harm others is more stringent than duties aimed at achieving other goods. Since the maximin principle is what the former duty requires under the circumstances, it has priority morally speaking over other valuable ends. Thus, while the end of fostering fraternity among citizens is certainly worthwhile, we should not violate the duty not to harm others in order to foster fraternity as we would if we asked the less well off to do with even less so as to achieve overall equality. Within a social framework governed by the maximin principle there will be room for these other values to operate. For example, people may give any income over and above the minimum to benefit those least well-off thereby making them better off yet. In fact, they may even have a duty to do so.³⁰ But the important point is that we cannot socially compel people to do so if society is in fact organized according to the maximin principle, since doing that would lead the least well off position to become less acceptable yet.³¹

Objection 5: Using Primary Social Goods is Unfair to the Handicapped.

The maximin principle will leave handicapped persons with too little, because with certain kinds of handicaps living a decent life requires a large number of resources above the minimal acceptable share under maximin, and perhaps even more than the maximum allowable share. If such people are to lead decent lives, they should get more resources.

Reply

Here we have reason to deviate from the maximin principle in the actual world. How far is not completely clear. While giving most people a relatively multi-purpose package of resources and liberties is to leave them unharmed this may not be so of all. Given that people with handicaps may well not be able to use the primary social goods they are allocated in very many ways, we do harm them when the package allocated is not sufficiently useful to leave them with a decent existence. So the not-harming constraint would entail that we allocate the handicapped a package of resources that they could in fact make decent use of. The underlying argument thus gives us reason to deviate from the maximin principle where real handicaps are involved.

Still, perhaps the maximin argument will not leave the handicapped with very much. If the handicaps in question are such as to be handicaps under any possible social situation, the argument so far will leave the handicapped entitled to resources allowing them fewer capabilities than everyone else. For, given that the baseline for making judgements of harm is determined as suggested above, people so handicapped will have a rather low baseline for comparison.³²

Here I think a public goods argument can come to the rescue. As a matter of lucky fact, most people do not in fact want to see the handicapped make do with just a bare minimum. Thus most all would be willing to pay something for the support of the handicapped so long as others did so as well. Taxation for the support of handicapped people at a reasonable level of well-being could be supported on these grounds, whether the not-harming rationale applies or not.³³

A Virtue of the Main Argument in Generating the Reply Strategies

These objections bring us to a virtue of the argument for the maximin principle. It is not normally a convincing answer to a hypothetical counter-example directed against a moral principle

that the postulated example is not likely to happen. If violations of a moral principle are intrinsically wrong it will count against a moral principle that its application would rule out something we do not think wrong. But the argument above for the maximin principle treats it as a means (a constitutive means in actual circumstances) for satisfying a more fundamental but perhaps overrideable moral principle, that of minimizing maximal harm. Where the maximin principle does not minimize such harm, we should deviate from it. This is what our replies regarding handicaps traded on, as did the restriction of the principle to cases of interaction. Furthermore, if we are non-absolutists about the minimax harm principle, we should give up the maximin principle just where the minimax harm principle is over-ridden by other considerations. This feature of the justification was used in replying to the first, utilitarian, objection. It is a virtue of the overall strategy of argument that it allows us to justify deviations from maximin, while leaving it intact for the broad range of cases it was meant to cover. The modifications are not *ad hoc*, because they are necessary to keep the use of the principle in line with its justification.

SOME CONCLUDING MORE ABSTRACT METHODOLOGICAL REMARKS

I'd like to conclude with some remarks about where this leaves my defense of the maximin principle. I began by providing a four-premise argument for the maximin principle, as well as auxiliary arguments on behalf of each of the premises. While these auxiliary arguments have some force, that force is not indubitable. In this last section I want to place the overall argument in perspective and argue that there are two sorts of reasons for taking it seriously, despite the less than certain nature of its premises. Both sets of reasons have to do with considerations of overall conceptual coherence. My first claim is that only some rival positions will be able to deny each of the premises of the argument. Hence only those will be able to take advantage of the cumulative

doubt stemming from multiple less than certain premises. My second claim is that discovering an underlying coherence between acceptance of the maximin principle, and acceptance of a duty not to harm can provide those with some antecedent commitment to both with additional reason to accept each one.

Why the Cumulative Effect of Dubitable Premises

Should Not Comfort Most Critics

None of the premises in my argument is indubitable. This should make us cautious but not entirely doubtful about accepting the conclusion on these grounds. The reasons we might have to doubt any given premise are not always compatible with reasons to doubt either the conclusion itself or with rationales for alternative positions. That limits the available positions that someone who rejects the argument might take. Let me explain.

Deciding on a theory of justice is a holistic endeavor, sensitive to arguments, but also dependent on the candidate theories that are available as alternatives. What we should believe depends partly on what positions are left when we have ruled out the alternatives. The fact that the argument contains a number of controversial premises should not comfort all opponents equally. Someone who is agnostic about what justice requires will find some support for that view in noting that the probability that the entire argument is correct is a function of the probability of each of its premises multiplied by the probability of the rest.

But things are not so simple for those with commitments to rival views about justice. Each of them requires that some particular premises in the argument be incorrect, on pain of undermining their own favored views. Even if consequentialists need neither affirm nor deny the last three premises of the argument, they must deny the first. Thus, the only premise whose uncertainty should

comfort them is the first. And, of course they do deny it, and they are not clearly wrong to do so. Still, their position is no more likely to be true than the first premise is likely to be false. Uncertainty over the remaining premises should not comfort them.

Furthermore, some of the most forceful arguments for the first principle, and against consequentialism involve examples in which aiming at overall best consequences would require the sacrifice of a few. The one form of consequentialism that is somewhat immune to this sort of objection builds a maximin principle³⁴ in to the assessment of consequences. This version would rank outcomes from worst to best by the situation of the least well off person in each outcome. The theory is somewhat immune to sacrificing-type objections because it never asks anyone to do without more than the worst off person in each of the available alternatives. So it has an answer of sorts to the objection that it unfairly asks some to sacrifice for others - namely that the alternatives ask for even more sacrifice from other people.

There are, I think, reasons to prefer nonconsequentialist theories to even this form of consequentialism. But my point here is just that a consequentialist theory will only be correct if the first premise of the argument is false, and that even then some form of maximin theory would still be viable. The best consequentialist defence against an egalitarian non-consequentialist position might be to accept a version of consequentialism that would endorse similar egalitarian conclusions.

The libertarian is typically in not much better position. While there is not just one particular premise that must be false if libertarians are correct, most libertarians will be precluded by their own commitments from rejecting the importance of the harming/not helping distinction.³⁵ They must find some way of arguing that I have mis-characterized the upshot of paying attention to that distinction so that it actually favors their position rather than maximin.³⁶ We could put the point the other way

around; as far as the advocates of libertarianism find their own position plausible, they should find many of the premises of the argument I offer plausible as well.

Rival nonconsequentialist egalitarian theories of justice may be in a somewhat better position, but it isn't obvious that they are. As non-consequentialists, they may share a commitment to the first premise of the argument but they need not. One can be a non-consequentialist for reasons besides accepting a duty not to harm others. If they share a commitment to that premise, once again they cannot take advantage of all of the uncertainty of each of the premises. If they don't share that commitment then perhaps they can. In any case, if some such rival theory is correct the maximin theory will not have gone too far wrong, since it is itself a rather egalitarian distributive theory.

The upshot of this is not, of course, that the argument itself is stronger than it is. Any uncertainty about the premises should ramify through to uncertainty about the conclusion. But, advocates of particular rival theories cannot argue that the chance that their own favored theory is correct is equal to the chance that my argument fails. For the argument may fail in a way that also does damage to the main arguments for their own theory. And some of those failures would actually make it more likely that egalitarianism is correct.

A Final Point About Coherence

This brings me to my second abstract point. Most of us accept the sorts of political views we accept partly as a result of arguments, but also partly because we find the arrangements described inherently attractive. Those of us who find maximin theories plausible do so out of some attraction to egalitarian theories and perhaps some sense that pareto optimal deviations from egalitarian distributions should not be objected to. If we also accept the nonconsequentialist conception of the duty not to harm, our justification for accepting each might actually increase once we discover that

there is a certain congruence of perspective when they both are applied to the distribution of social resources. And if we find further that the proffered rationale explains why what might seem to be plausible counter-examples to the theory leads us to want to provide for exceptions to the requirements of a pure maximin theory, that should increase our degree of justification to a still greater degree. We wind up in a position to see why these qualifications are not ad hoc modifications to a view the spirit of which is at odds with the modifications, but rather an adjustment of the way that theory has been applied to better capture that spirit.

Notes

1. New argument or not, my debts to Rawls in this paper are obvious and enormous. While I don't think his official argument is the one I'm presenting, many of the things he says in *Theory of Justice* and subsequent articles have been incorporated into my defense of the main argument of this paper. My argument also owes a lot to work by G. A. Cohen; it was suggested to me by reading his criticisms of Nozick's treatment of original acquisition and the key idea, that one can harm by removing opportunities is from that discussion. One might even wish to think of this as Cohen's argument for the maximin principle, except that he would dissent from the conclusion and several of the premises. I also owe debts and thanks to many people who took the time to comment on earlier drafts of my paper. [Rest of note omitted to preserve anonymity.]

2. Rawls sometimes reserves the term 'maximin' for what is required by the difference principle. But Rawls's special conception of justice (including the difference principle) is itself defended as justifiable by a kind of maximin reasoning. I think it most plausible to read Rawls as presenting this choice as an application of his general conception of justice (requiring that all primary social goods be distributed so that any inequalities are to the advantage of the least well off) to our own favorable situation. His overall theory then would require maximizing the minimum share of all primary social goods. See, John Rawls, *Theory of Justice* (Cambridge; Harvard, 1971) pp. 62, 151-3 and 302-3.

Some readers may wish to resist attributing this idea to Rawls on the grounds that we have no way to measure shares of primary goods besides income and wealth. I don't think this is much of a problem so long as we can roughly rank packages of primary social goods. Rawls's own discussion of individuals choosing the special conception as the least risky option behind a veil of ignorance

indicates that Rawls too thinks this sort of thing can be done. But, whatever Rawls's own views on the matter, when I speak of the maximin principle in this paper, I will be referring to a principle for the distribution of all primary social goods.

3. By a stringent duty, I mean a duty that is not easily overridden by other considerations or duties. Since I think that the duty not to harm prevents us from seriously harming some even to prevent similar harms from coming to a greater number of others, I think it is stringent in this sense. This is compatible with its being overridden where the harms prevented would be very much greater than the harms caused so long as it is not overridden where the harms are only somewhat greater.

4. In other words, we should choose the option in which the person we most harm is harmed less than we would harm the person we most harm were we to choose an alternative option.

5. For work broadly in this traditions see: Robert Audi, "Intuitionism, Pluralism and the Foundations of Ethics," in W. Sinnott-Armstrong and M Timmons (eds.) *Moral Knowledge?* (Oxford: Oxford University Press, 1995) 1-136, Berys Gaut, "Moral Pluralism", *Philosophical Papers*, 22 (1993) 17-40, David McNaughton, "An Unconnected Heap of Duties?", *Philosophical Quarterly* 46 (1996) 433-447, W. D. Ross, *The Right and the Good*, (Oxford: Clarendon Press), Bernard Williams, "What Does Intuitionism Imply?" in *Making Sense of Humanity*, (Cambridge: Cambridge University Press, 1995) 182-191.

6. Philosophical defenses of this idea can be found in Warren Quinn's "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing" *Philosophical Review* 98 (1989) pp. 287-312, Philippa Foot, "Killing and Letting Die," in Garfield and Hennessey, eds., *Abortion and Legal Perspectives*, (Amherst: Univ. Of Massachusetts Press, 1984), Robert Nozick, *Anarchy, State, and Utopia*, (New York; Basic Books, 1974) pp. 26-53, and F. M. Kamm, *Morality Mortality, Volume II, Rights, Duties and Status*, (New York; Oxford University Press, 1996).

7. I think that it is permissible to harm one person to avoid great harm to many others, at least where the harm to the person is relatively small, or where the I am harming someone to prevent catastrophic harms to many others. (See Chapter 19 of Michael Walzer's *Just and Unjust Wars*.)

8. This objection was pressed forcefully by [omitted], drawing on a way of defining harms made by Joel Feinberg in Harm to Others, (Oxford; Oxford University Press, 1984) 31-36.

9. We should note that the issue here is not about the relevance or coherence of the distinction between the kinds of agency that underlies the harming/not helping distinction. The moralized duty not to harm requires such a distinction, plus an account of what people are entitled to. The moralized version claims that the duty rules out acting in certain ways to deprive a person of what they have a right to, whereas the nonmoralized version requires only that the duty rules out acting in certain ways to make someone worse off than they would be without our acting in that way.

10. Philippa Foot “The Problem of Abortion and the Doctrine of Double Effect,” *Oxford Review*, No. 5 (1967) reprinted in *Virtues and Vices*, (Berkeley; University of California Press, 1978) pp. 19-32, and Judith Thomson, “Killing, Letting Die, and the Trolley Problem,” *The Monist* (1976).

11. Ross himself seemed to think that when such conflicts arose there was never a further principle to guide the choice about what to do. This claim is at least not obviously correct, even if a complete set of principles for resolving such conflicts is not to be had. Even if for many sorts of conflict there is not much to say ahead of time about how they should be resolved, it is surely possible that in some cases there might be quite a lot to say. Where the actions in question are the actions of groups, it is even more plausible that principles will be needed to provide guidance and coordination. Thus I don't think we move beyond Ross in an implausible way if there are any principles to guide us when the duty not to harm conflicts with itself or other prima facie duties, at least so long as we don't say there must always be such principles.

For example, if our measure of individual well-being is a one which gives greater priority to securing a certain minimum for a person.

13. Alasdair Norcross would not agree. See, "Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches," *Analysis* (1998) 58(2), 152-158.

14. This broadly Kantian idea is developed by Robert Nozick in *Anarchy State and Utopia*, (New York; Basic Books, 1974) 30-42 and by Warren Quinn in "Actions, Intentions and Consequences," *Philosophy and Public Affairs* (1988), reprinted in his *Morality and Action* (Cambridge; Cambridge University Press, 1993) 175-193.

15. This section of the paper is especially indebted to G. A. Cohen's discussion of Robert Nozick's treatment of the "Lockean Proviso," requiring that original appropriation leave "enough and as good" for others, cited more fully below.

16. Perhaps because in that world different talents were considered socially valuable, or perhaps because a certain amount of money were distributed by lottery to a lucky few.

17. Cohen makes good use of this idea to criticize Nozick's justification of capitalist property rights in his paper, "Self-ownership, World Ownership and Equality," in Frank Lucash, ed., *Justice and Equality Here and Now* (Ithaca; Cornell University Press, 1986) incorporated into chapter 3 of *Self-Ownership, Freedom and Equality*, (Cambridge; Cambridge University Press, 1995). Cohen's discussion there is focused on what it is to make a person worse off, rather than explicitly discussing what can count as a harm, but I think that his point is related. We can see Nozick as trying to justify the original appropriation by showing that it harms no one, since it "makes no one worse off." Cohen's criticism shows that appropriation of nature carrying with it full capitalist property rights does in fact make people worse off, and hence harms them.

18. The point is implicit in Cohen's "Self-Ownership, World-Ownership and Equality," and explicit

in Jeremy Waldon's "Welfare and the Images of Charity," *The Philosophical Quarterly*, 36 (1986), pp. 463-82, reprinted in *Liberal Rights: Collected Papers 1981-1991*, (Cambridge; Cambridge University Press, 1993) pp. 225-249.

19. If you think you were that well off all along, think of the thought experiment like this. First you buy the ticket for the lottery at a slightly better than fair price, say five hundred dollars. Then the events which determine the outcome transpire and you wind up with four hundred dollars. It doesn't seem that buying the lottery ticket at a better than fair price made you worse off. But you are worse off when it is all over. So it seems that the lottery itself made you worse off.

20. To show that the likelihoods seem relevant we can adjust the probabilities of the different payoffs.

21. Things are actually a bit more complicated than that, because while the choice of arrangement partly determines a person's well-being, the person's responses to being in that arrangement, as well as luck, will also partly determine their well being. Consistent with the non-consequentialist cast of this argument, the component depending on the agent's own actions will not be morally relevant to determining whether we have harmed her by putting her in that situation. Hence, it might be acceptable to use her well-being if she acted to further her interest in that situation, as the expectation for each alternative.

22. It seems obvious to me that we should not include outcomes that result from such gratuitous worsenings of a person's situation if figuring the baseline for harming, but I'm not sure I can provide a non-question-begging rationale. It would be odd to allow them into my general way of thinking about harms, since it would seem to make whether I harm you or not by one of my actions partly a function of the other even worse things I might possibly do to you instead. In any case, it may not matter since they will mostly have a low antecedent probability insofar as no one would have any

good reason to aim at them, and at least one person would have a strong motive to avoid them. Thus, if the baseline is a function of the probable possible outcomes, excluding these does not much effect where it lies.

23. I am not here advancing expected utility as the correct way of measuring prospects, though I think it will have some appeal to those with utilitarian sympathies. I am using this to illustrate the general idea that an appropriate baseline should be a function of all the alternatives possible before human agency is factored in. We could use another measure of a person's prospects, and yet still be working with the same basic idea.

24. I think that this is part of common sense, but exactly how to work the idea in to a philosophical account of the doing/allowing distinction will be tricky. For some of what a person does to themselves may be possible only because of what we have done to or for them. So perhaps these things are doings of ours insofar as they result from our actions. I'm inclined to say that where our action leaves the person a real choice about what to do, the outcome is not something that counts as a doing of ours. But this can only be a first approximation, or we will have problems counting the upshot of jointly undertaken courses of action as anybody's doing. Our intentions will likely play a role in the worked out version of the distinction.

25. Rawls, in answer to objections that the use of primary social goods for purposes of comparison is unfair to the handicapped remarks that he thinks that it may be best to use some other mode of comparison (perhaps Sen's capabilities) for such individuals, but that primary social goods are still useful as a means of comparison in ordinary cases. *Political Liberalism*, (New York; Columbia University Press, 1993) p. 185.

26. I've heard Joshua Cohen make a similar reply to utilitarian objections to Rawls.

27. Nozick, *Anarchy State and Utopia*, (New York; Basic Books, 1974) 190-97.

28. Nozick, p. 185. Barry, *Theories of Justice*, (Berkeley; University of California Press, 1989) pp. 234-41.

29. The argument is suggested by G. A. Cohen's, "Incentives, Inequality and Community," in *The Tanner Lectures on Human Values* (Salt Lake City; University of Utah Press, 1992) pp. 261-330.

30. Such a duty might in fact be derived by a similar argument to the one that generated the maximin principle. Or it may just be a duty of fraternity or beneficence.

31. We can of course compel them to give whatever is necessary to bring compliance to the maximin point. But, if we demand more, we are harming those at the bottom.

32. This way of judging the baseline will be appropriate only for naturally occurring handicaps. For those caused by human beings, the people causing the handicap itself will be responsible for the harm done, and self-inflicted handicaps would be the responsibility of those who bear them. It would likely be reasonable for citizens to use public moneys to provide for people in both categories nonetheless, by establishing an insurance pool to protect us all from either self-inflicted harm, or liability for harm to others. The appropriate level of compensation might be determined in the manner suggested by Allan Gibbard, in "Health care and the Prospective Pareto Principle," *Ethics* 94, (Jan. 84) 261-282.

33. A not-harming rationale can be constructed to allow taxation for public goods. Forbidding such taxation would be depriving people (including the least well off) of a means of improving their situation by cooperating with others. Balanced against this is the harm of taxing those who do not share in the goals motivating the tax. But if most people have goals that would be furthered by taxation for those ends, it can be in the interests of all that taxation for public goods be allowed, even when a majority would not benefit from the various taxes individually. This line of argument would also provide a ready answer to critics of maximin who argue that it unrealistically rules out taxation

to fund the arts or humanities. But that is an argument for another day.

34. See, Joseph Mendola “An Ordinal Modification of Classical Utilitarianism,” *Erkenntnis* 33 (July 1990) pp.73-88, and Derek Parfit “Equality or Priority?” The Lindley Lecture (University of Kansas, 1995) pp. 19 ff.

35. For example, Milton Friedman and Robert Nozick. Friedman's argument relies at various points on the idea that my freedom to act is limited by your freedom from interference. See *Capitalism and Freedom* (Chicago; University of Chicago Press, 1962) pp. 12-15 & 25-26. Nozick's defense of original appropriation of property so long as it is limited by a Lockean proviso that others not be made worse off, can be seen as an argument that such appropriation is justified because it harms no one.

36. Here I find G. A. Cohen's arguments cited above relevant and entirely persuasive.