

Intuitive Maximin

JOSEPH MENDOLA
University of Nebraska
Lincoln, NE 68588-0321
USA

One standard objection to familiar utilitarian consequentialism queries its troubling commitment to the maximization of overall value irrespective of distribution, for instance among the well and badly off. Call this 'the objection from distribution.'

The simplest and most obvious alternative form of consequentialism deploys some sort of maximin principle. Maximin principles maximize the well-being of the worst off. *Lexical* maximin rules in particular, which are perhaps the simplest and most obvious subtype, maximize first the well-being of the worst-off, and then in case of ties among the worst-off, maximize the well-being of the second worst-off, and so forth. Maximin principles provide an obvious route to the unification of plausible concerns with maximization and with distributional equity.

And yet it has seemed that maximin principles are not really a plausibly intuitive way to dispose of the objection from distribution. Almost no one believes that *direct* application of such principles is remotely plausible. John Rawls was the most prominent proponent of maximin principles.¹ But even he cushioned the effects of maximin in certain ways which are not available to maximin consequentialism. In Rawls, maximin principles only govern life expectations for primary social goods of representative members of basic social groups in just societies, and indeed merely in regard to economic expectations inside a fixed context of political liberties and equality of opportunity. Rawls, himself, sug-

1 John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press 1971)

gested that maximin principles will be in conflict with common sense when applied in a more straightforward way.²

I believe that maximin consequentialism would be, in fact, found intuitive by properly informed and reflective common sense. But my point here is smaller: there are certain concrete cases which seem intuitively to favor maximin, at least against familiar competitors, and certain concrete cases which seem intuitively to count against it. Yet any legitimate difference between those cases must rest on the deep moral significance of the difference between acting and refraining from action, or to be more exact, between failing to aid and harming. But the moral significance of that distinction is itself quite controversial. It is not beyond commonsense dispute, and is, indeed, one of the central outstanding issues between consequentialism and its deontological and other critics.

In its simplest forms, consequentialism does not support a moral difference between failing to aid and actively harming when they have the same consequences. Even in sophisticated forms, consequentialism often does not support a deep such difference. Many, of course, consider this yet another objection to consequentialism.

But my goal here is merely to link two standard objections to standard forms of consequentialism, and support a conditional claim: If consequentialism can legitimately undercut the apparently intuitive objection that it does not underwrite a sufficiently deep difference between failing to aid and harming, then it can evade the objection from distribution by deploying a maximin principle.

Honesty requires that slight complication about depth. I believe that the difference between acting and refraining doesn't have the deep moral significance suggested by some elements of common sense. Nor does that between harming and not aiding. But I believe that there is some difference, just not a very significant one.³ I believe that these differences aren't significant enough to support a differential treatment of the cases we will face here. Because there is no sufficient moral difference between the intuitive cases which seem to suggest maximin and those which seem initially to intuitively count against it, we are forced to choose between our intuitions, between those which favor maximin and those which count against it. But I may be wrong that the complication about depth

2 Rawls, *A Theory of Justice*, 153-7

3 Joseph Mendola, 'Consequences, Group Acts, and Trolleys,' *Pacific Philosophical Quarterly* 86 (2005) 64-87

is relevant to reality, and in any case the difference introduced by that complication seems unimportant to the argument we will traverse here.

Besides denying a deep distinction between not aiding and harming, I also believe that the very facts which show that there is not the requisite difference between the cases favor the intuitions which suggest maximin and undercut the intuitions which apparently constitute objections to it. But I will not attempt to establish those large claims here. My job here is merely to probe the interaction of, on one hand, the intuitive status of maximin and, on the other, the significance of the difference between harming and failing to aid. Maximin will be seen here to be suitably intuitive *if* that distinction is not of sufficient normative significance.

I

Our first case seems to favor maximin, at least if such a case can tell us anything about how to evaluate alternative outcomes. Consider the intuitive force of Ivan's claim in *The Brothers Karamazov* that he, unlike God — that he *against* God — would not consent to the torturing of one child to create great happiness for the rest of mankind. The suffering of that consequently worst-off individual intuitively trumps a concern for the well-being of the rest. Ivan says to his brother:

'[I]magine that you yourself are building the edifice of human destiny with the object of making people happy in the finale, of giving them peace and rest at last, but for that you must inevitably and unavoidably torture just one tiny creature, [a] child who was beating her chest with her little fist, and raise your edifice on the foundation of her unrequited tears — would you agree to be the architect on such conditions? Tell me the truth.'

'No, I would not agree,' Alyosha said softly.⁴

Many share Ivan's and even Alyosha's intuition that there would be something morally and humanly reprehensible in such an act. Of course, the child isn't just worst-off. The case involves torture; indeed, the torture of an innocent. So many ethicists will claim that this case doesn't reveal anything about the relative evaluation of outcomes. The outcome Ivan considers may be a better one, they think, but it is just not appropriate to get there by torturing someone. Part of my job must be to show that the torture isn't a distracting and misleading feature of the case, or, rather, that it isn't a distracting and misleading feature of the case if there

⁴ Fyodor Dostoevsky, *The Brothers Karamazov*, Richard Pevear and Larissa Volokhonsky, trans. (New York: Farrar, Straus and Giroux 1990), 245

isn't a deep distinction between harming and failing to aid. Please suspend for the moment your belief that I cannot do that.

Now consider cases, in certain ways closely analogous cases, which seem to pull against maximin. The most thorough negative discussion of maximin is Larry Temkin's. Like Rawls', Temkin's discussion is specifically focused on the relative justice of circumstances, as opposed to other virtues. But this seems a detail that we can properly ignore here.

Temkin begins with a situation S with one well-off group, one not so well-off group, and one still less well-off individual. Here are Temkin's main contrasting cases and his main conclusion rooted in intuitive reactions to those cases:

A, B, and C are alternatives to S. In A, the worst-off person would remain unchanged, the better-off group would be slightly lowered, and everyone else would be dramatically raised. In B, the worst-off person would be slightly raised, but everyone else would be lowered, and while the others who initially fared poorly would still not be as badly off as the worst-off person, each of them would lose more than the worst-off person gained. In C, the worst-off person would be raised slightly more than in B, the better-off people would be raised significantly, and the others would be lowered to the worst-off person's new level, each (again) losing more than the worst-off person gained.

On a Maximin Principle of justice focusing on the worst-off person, B and C would both be more just than A, with C most just. Many find this unacceptable.⁵

I admit the force of these intuitions. Indeed, let me be yet more concessive. Temkin's analysis not only plausibly captures our initial intuitions about these cases, but intuitive reasons for these reactions. Maximin is intuitively incorrect for two closely-related reasons, according to Temkin. Here is the first:

Although considerations of justice may focus our attention on the worst-off person, and we may even be *most* concerned about her situation, surely her plight is not our *only* (significant) concern regarding justice.... [I]t is implausible that we should be deeply and genuinely concerned about benefitting the worst-off person, but not concerned at all about benefitting those who fare very poorly but are (ever-so-slightly) better off than she. In bringing about B or C rather than A, members of S would be directly harming, as well as failing to *significantly* benefit, many who are themselves quite badly off — all for the sake of slightly benefitting the worst-off person.⁶

5 Larry S. Temkin, *Inequality* (Oxford: Oxford University Press 1993), 103

6 Temkin, *Inequality*, 103

Here is Temkin's second reason:

Let A^* and C^* be large populations.... Let A^* be perfectly equal with everyone faring very well, C^* very unequal with the worst-off group faring very poorly. Finally, assume that in Rawlsian terms A^* is *perfectly* just, and C^* terribly unjust. Clearly, the committed judgment of Rawlsians ... would be that regarding justice A^* would be much better than C^* . But notice, A^* and C^* might be represented by A and C ... except for the worst-off person. Yet ... focusing on the worst-off person, Maximin would rank A and C the exact reverse of A^* and C^* . This seems implausible.... [F]ocusing on the worst-off person ... may seriously distort our judgment because the worst-off person's condition may not accurately reflect the situation's overall justness.... [It] may simply be an anomaly or fluke.⁷

So where are we? Temkin's cases and the general rationales he cites have some genuine intuitive force. But we also face Ivan Karamazov's in some ways very similar case, which seems to pull our intuition in the other direction. We must reconcile these apparently conflicting intuitions.

II

So our question becomes: What is the difference between Temkin's cases and Ivan's? Temkin himself notes what I take to be the key *surface* difference between the cases:

Although some might think it would be unjust to *bring about* A if one *started* in B or C , many believe A would be more just than B or C , and that starting in S justice dictates bringing about A rather than B or C .⁸

In other words, it apparently matters what situation we start with, because it apparently matters whether a situation is brought about by one's action or merely remains the case because of one's inaction. It apparently matters whether someone's horrible situation is a result of active torture by Ivan or is merely something we ignore in order to favor the interests of better-off others. It also apparently matters whether we will have to actively harm those others so as to merely aid the truly worst off. The intuitive difference between the cases rests, as I suggested earlier, on the apparently intuitive difference between harming and failing to aid, indeed perhaps on more than one such difference.

⁷ Temkin, *Inequality*, 103-5

⁸ Temkin, *Inequality*, 103

Temkin's discussion is accompanied by relevant footnotes which reinforce this point:

If we learned that A's situation resulted from the (concealed) presence of a tortured slave, whose tormented struggles enabled everyone else to fare so well, we might completely revamp our judgment regarding A's justness.⁹

So in effect he grants the intuitive force of Ivan's case. And he says this about the difference:

The relevance of the starting point to our judgments may reflect two important strands of moral thinking. The first corresponds roughly to a distinction like that between the good and the right to which many nonconsequentialists adhere. The good or just situation cannot always be rightly or justly brought about. The second corresponds roughly to the asymmetry many see between harming and not helping. Although it may seem unjust to harm those already worse off than everyone else so as to benefit others, it may not seem (as) unjust not to help those already worse off than everyone else so as to benefit others.¹⁰

Temkin is certainly no friend of maximin. His work is its most thorough criticism. But the friends of consequentialist maximin should accept that this part of his analysis of what is going on in these cases is essentially correct. The intuitions about these various cases depend directly for their probity on the presumed deep significance of the difference between harming and not helping, or analogous presumptions that otherwise better situations cannot be justly brought about. The friends and enemies of maximin should agree so far.

But the rub is this. Many consequentialists and at least some strands in common sense do not accept that these factors have the deep significance which Temkin presumes. We certainly have no normative consensus that the paradigmatic consequentialist view of these issues, the view that these alleged differences are not deep differences, is incorrect. If consequentialists can properly insist that there is no deep difference between harming and failing to aid, then the intuitive case against maximin is not so obvious.

Certainly if the moral difference between aiding and not harming is deep, then maximin is intuitively implausible. I grant that. But what if it is not? If it is not legitimate to treat Temkin's and Ivan's cases differently on such a ground, then what are we to conclude? The worst that could reasonably be said against maximin on that presumption is that com-

⁹ Temkin, *Inequality*, 105

¹⁰ Temkin, *Inequality*, 104

nonsense intuition cannot resolve its plausibility, I believe. But in fact on such a presumption we can do better than that.

If, for simplicity's sake, we presume that there is *no* morally significant difference between acting and refraining, and between harming and not aiding, that would not intuitively serve to *remove* moral responsibility for the effects of one's action, but rather to extend it to effects of one's non-action. Consequentialists characteristically think we are morally responsible for more that happens. And it is one's responsibility for the suffering of the child in Ivan's case which so deeply evokes our intuitive concern.

Deontologists think that torturing someone is specially wrong in a way that letting someone be tortured is not, and some ethicists think that it is worse if someone is tortured than if a tree falls on them and does comparable harm. But consequentialists characteristically argue that letting someone be tortured is as bad as torturing, and that letting a tree fall on someone may be as bad as letting them be tortured, *not* that torturing is as normatively insignificant as some deontologists presume letting someone be tortured to be, *not* that letting someone be tortured is as normatively insignificant as some believe letting a tree fall on someone to be. That indeed is the intuitive implication of a collapse of the customary distinction between harming and not aiding, to *extend* responsibility, in other words to undergird our normative response to Ivan's case and suggest a generalization, and to undercut our normative response to Temkin's cases. It is not the torture in Ivan's case which misleads our intuition, but rather the lack of responsibility presumed in Temkin's cases, and indeed generally presumed by unreflective common sense when it confronts such cases.

III

Since this is the key point, we should pause to consider natural resistance. Temkin does at least sometimes suggest a situation of choice among the options he provides, and talks of harms from and gains from a given baseline, and that may invoke responsibility of at least a weak sort. But his goal is proper intuitions about the justice of outcomes, independent of whether or not anyone has responsibility for them. And we may think that the way to uncover these intuitions is to remove as much responsibility from the situations to be considered as possible, not to dramatically heighten it in the way that torture does. In other words, we may think that Temkin's procedure is basically right, but that he should purify his cases yet further to remove those elements which may even weakly suggest responsibility for the outcomes in question.

But consequentialists should disagree. They should insist that our real and unconfused evaluation of outcomes is revealed when the full weight of our responsibility for anything which happens because of our choice is felt.

Of course, things are probably not quite that simple and clean. There are complications. First of all, there is a series of differences between Temkin's cases and Ivan's torture which I mostly lumped together in the previous paragraphs. There is the difference between torturing and letting someone be tortured by someone else. And then there is the difference between someone's being tortured and otherwise harmed in the same way. And then there is the difference between their being harmed and their starting out in a bad situation from which they might be rescued. Moreover, there is our standing complication. There can be *some* moral difference between harming and not aiding with the same consequences even if there is not a deep and significant difference.

As I said, I believe that the facts in the end are that there is *some* difference between harming and not aiding, and indeed that there is some difference along each of the steps we might distinguish between Temkin's cases and Ivan's. But these are not sufficient differences, nor indeed the right kind of differences, to constitute the torture present in Ivan's case to be intuitively misleading while the lack of responsibility in Temkin's cases is not. It is still rather the presumed lack of responsibility in Temkin's cases which misleads our intuition in the crucial way.

And my own analysis is not idiosyncratic for a consequentialist. Paradigmatic consequentialists who grant some difference between the cases will still attempt to extend our responsibility for outcomes. According to this whole class of views, not the many small but real moral differences between torturing someone and leaving them in an equally painful state, but rather the general false presumption that we would not be responsible in Temkin-style cases, is what dominantly drives our false intuition of a normative difference between Temkin's cases and Ivan's. We are responsible for the effects of our actions and inactions in the most morally relevant sense, consequentialists think. And commonsense intuition, if properly purged of the errors underlying the belief that there is a deep moral difference between harming and not aiding which eliminates our responsibility for the effects of not aiding, an alleged deep moral difference which does not in fact exist, would not conflict with maximin. My own detailed beliefs about the small differences between harming and not aiding have no probative significance here. But the general tendency of consequentialists to extend responsibility, which has some support in common sense intuition, suggests that even if there is some morally relevant difference between Temkin's and Ivan's cases, it will be relatively minor, and that it will indeed be such as to undercut

our intuitions about Temkin's cases and support out intuitions about Ivan's.

With the complications or without, this is a very controversial feature of consequentialism. Deontologists and others will object. They do not admit an extension of responsibility beyond the surface commitments of common sense. But here I rest merely with the conditional claim, that if there is no deep difference between harming and failing to aid, then maximin is suitably intuitive. If this is right, then what have seemed two distinct objections to standard forms of consequentialism—the objection from distribution and deontology's objection to its characteristic extension of responsibility—are connected in an unexpected way.

IV

There is one further and important complexity which must be addressed. Consider Temkin's first rationale against maximin, which stresses that maximin would somewhat sacrifice the well-being of many pretty badly-off persons to benefit the very worst off. We haven't directly discussed that objection. And his second rationale against maximin reinforces the worry that it is inappropriate to focus like maximin merely on the worst-off individual. We haven't yet considered these central elements of Temkin's alternative explanation of what is going on in these cases.

But note that if we prescind from the elements of Temkin's cases which evoke questionable intuitions about the normative significance of the difference between harming and not aiding, and assume that the agent in question has moral responsibility for everything in the various options, some recent and quite intuitively sensitive deontological accounts suggest a treatment in accord with maximin consequentialism.

For instance, Scanlon has granted that 'in situations in which aid is required and in which one must choose between aiding a larger or smaller number of people all of whom face harms of comparable moral importance, one must aid the larger number. On the other hand, [the correct view] does not require, or even permit, one to save a larger number of people from minor harms rather than a smaller number who face much more serious injuries.'¹¹ In a situation where one is responsible for saving large numbers from small harms or a few from larger harms, save the few. That may suggest we should generally focus on helping

11 Thomas M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press 1998), 238

the few worst off people at the expense of larger numbers of badly off people.

But, as we might expect, there are complications. Perhaps the course Scanlon suggests will still seem plausible even if those facing the larger injuries will end up all-told better off one way or another, whether they are injured or not. Still, at the very least that makes our intuitions less clear.

A second complication is that maximin requires other tough conclusions. Imagine that a million people are going to be tortured for fifty years, and one person for fifty years and a day. Maximin suggests that it would be better to reduce the torture of that one person by a day, rather than eliminate everyone else's torture completely, and this may seem quite implausible. But note that in cases of that sort which are intuitively clearest, it is correspondingly unclear whether the one day out of fifty years would make any difference whatsoever to the person about to be tortured, who may seem worst off in only an abstract and mathematical sense. When the difference to the worst off is obviously significant, then individual intuitions may well differ. And of course to focus solely on cases of saving people from harms and tortures may somewhat mislead common sense in the ways we have already discussed. If we are not to be misled, the consequentialist may plausibly insist, perhaps we should consider a case in which we ourselves would have to provide the extra day of torture to the one victim to relieve the sufferings of the others. Deontologists should think twice before embracing that course of action, even in a conditional way.

Of course, there are other cases. Two otherwise identical individuals A and B are struggling to obtain some scarce medicine, which would enable A to live to a ripe old age rather than die peacefully now, and B to extend their life by a month rather than die in severe pain now. Maximin suggests that it would be better for B to get the medicine, while it may seem intuitively better for A to get it, since then it will do more overall good. Still, if we focus on a situation in which you, as a third party, would have transparently full responsibility for both these outcomes, in which you would have to torture B to death now so that the medicine would extend A's life for so long, we can see that there are intuitions which pull in the other direction.

I do not claim that maximin can make us happy in the face of all hard cases and tragic choices. No view can do that without illusion, at least if we are responsible for all the effects of our choices. And we have no consensus that we are not. I do not even claim that Ivan's case, or other similar cases, suffice to show that a maximin principle is to be favored over other relatively close and hence relatively unfamiliar competitors, say a sufficientarian principle which requires the minimization of the number of individuals who fall below some critical threshold of well-be-

ing. Rather, my point is this. Temkin provides the most thorough and careful negative treatment of maximin we possess. And while many believe that they know maximin is false, they know no such thing.¹²

Received: June 2003

Revised: September 2004

12 Thanks to three anonymous referees for very helpful comments.

