



Fairness

Author(s): John Broome

Source: *Proceedings of the Aristotelian Society*, New Series, Vol. 91 (1990 - 1991), pp. 87-101

Published by: Oxford University Press on behalf of The Aristotelian Society

Stable URL: <http://www.jstor.org/stable/4545128>

Accessed: 01-06-2018 20:30 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://about.jstor.org/terms>



JSTOR

The Aristotelian Society, Oxford University Press are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society*

V*—FAIRNESS

by John Broome

This paper presents a theory about fairness, as it applies to the distribution of goods between people. I shall concentrate particularly on random lotteries. Sometimes a lottery is the fairest way of distributing a good, and my theory explains, better than any other theory I know, why this is so. That is the main evidence I offer for it. But the theory is not limited to lotteries; it is intended to apply whenever goods are distributed between people. I shall use the fairness of lotteries as a guide to fairness in general.¹

I

It often happens that there are several candidates to receive a good, but the good cannot be divided up to go round them all. The good may be very important; it may even amount to the saving of the candidate's life. For instance, not enough kidneys are available for everyone who needs one. As a result, some people are denied treatment for their kidney failure, and consequently die.

For each candidate, there will be reasons why she should have the good, or some of it. (I mean *prima facie* reasons, which may be defeated by other reasons.) Amongst them will be the benefits, to the candidate herself and to other people, that will result from this candidate's receiving the good. When the good is the saving of life, these benefits will depend on how much the candidate enjoys her life, what responsibilities she has to other people, and so on. Then there may also be other reasons. One may be desert: some of the candidates may deserve the good, perhaps because of services they have performed in the past. For the moment, suppose all these reasons can be weighed against each other. (I shall question this later.) Then for some candidates, the reasons

* Meeting of the Aristotelian Society, held at the Senior Common Room, Birkbeck College, London on Monday 3rd December, 1990 at 8.15 p.m.

¹ Elsewhere, I have applied the theory to the distribution of divisible goods such as income. See my 'What's the good of equality?' in John Hey (ed.), *Current Issues in Microeconomics*, Macmillan, 1989, pp. 236-62, and my *Weighing Goods*, Blackwell, 1991, Chapter 9.

why they should have the good will be stronger, on balance, than for others. Let us call one person a 'better' candidate than another if she has stronger reasons in her favour.

How should it be decided which of the candidates should get the good? Several procedures might be used. One is to have some authority judge the merits of the candidates, and select the best. But this procedure has its costs. The job of assembling and assessing the necessary information may be expensive and time consuming. The responsibility of deciding who is to live and who to die (if that is in question) may be an intolerable emotional burden. Furthermore, the authority may not actually succeed in picking the best candidates. It may choose the candidates who best meet corrupt or prejudiced criteria, rather than the ones who are actually the best.

One procedure that avoids the costs and dangers of deliberate selection by an authority is to apply some fixed rule. (There is a risk of corruption or prejudice in setting up the rule, but not once it is set up.) And it may be possible to devise a rule that goes some way towards selecting the best candidates. For life saving, the rule of picking the youngest will do this. Age will certainly be one of the factors that helps determine which candidates are the best. Other things being equal, it is better to save a younger person than an older, because it does more good to the person who is saved: it gives her, on average, more years of life. So there is some correlation between a person's youth and how good a candidate she is.

A lottery is another procedure that avoids the costs and dangers of deliberate selection. Unlike a well-chosen fixed rule, though, it is no more likely to pick the best candidates than any others. So what advantage can it possibly have over a fixed rule? Plainly, only that it is sometimes fairer. But how can this be so? How can a lottery be fairer than a rule such as picking the youngest, which has a tendency to select the better candidates? Answering this question is the main test that has to be passed by any account of the fairness of lotteries. To answer it properly demands a particular theory of fairness in general; only this theory, which I shall describe in Sections III and IV, is able to explain adequately the fairness of a lottery.

II

Before coming to the theory, I shall first set out what I think it needs to account for: the facts about the fairness of lotteries that need to be explained.

First: a lottery is by no means always fair. It would not, for instance, be a fair way of choosing whom to award the prize in a violin competition. So in explaining the fairness of lotteries we shall need a criterion for distinguishing when lotteries are fair from when they are not.

Second: our account of the fairness of lotteries cannot simply be that lotteries are good tie breakers, appropriate only when the reasons in favour of different candidates are exactly balanced. The two examples below show that it is sometimes right to hold a lottery even when reasons are not exactly balanced. In any case, if a lottery were appropriate only for breaking a tie, its value would be insignificant. It will hardly ever happen in practice that reasons balance exactly. And if ever they do, the slightest change in one of them would mean they were no longer balanced. Then, if it was right to hold a lottery only for breaking a tie, it would no longer be right to hold one. So the value of a lottery would be lexicographically dominated by other values. (Section VI, however, qualifies this point.)

Furthermore, to say that lotteries are good tie breakers fails to explain their fairness. When there is a tie, it does not matter which candidate is chosen. What is required is simply a means of getting the decision made. A lottery is a handy means, even when no issue of fairness arises. When I cannot decide between two restaurants for dinner, I may toss a coin. This is not in order to be fair to the restaurants, but simply to avoid the fate of Buridan's Ass. When it comes to a choice, not between restaurants, but between candidates for some good, a lottery is sometimes more than just a handy means of getting the decision made when there is a tie. It is sometimes a better means than others because it is fairer. We, therefore, need a separate explanation of why it is fairer.

Thirdly: the fairness of a lottery does not consist solely in the fact that it overcomes the costs and dangers of deliberate selection by an authority. I have already explained that selection by a fixed rule is likely to be a better way of doing that. And even

when it is possible to choose the best candidates deliberately, without cost and without corruption or prejudice, there is still sometimes a case for a lottery. The following two examples make this point, and also the second point mentioned above.

The first example is about games. Most games begin by holding a lottery to settle which player starts in the most favourable position (playing white, say). Fairness requires this. But normally some players will be better candidates for the favourable position than others, as I defined 'better candidate' in Section I. For instance, usually more joy will be caused in total by the victory of one player rather than another, so a greater expectation of benefit would result from giving that player a favourable start. Let us suppose there is a referee who, without prejudice or corruption, is easily able to pick out the best candidate. It would still be wrong to leave the decision to the referee rather than a lottery.

The second example is a dangerous mission. Someone has to be sent on a mission that is so dangerous she will probably be killed. The people available are similar in all respects, except that one has special talents that make her more likely than others to carry out the mission well (but no more likely to survive). This fact is recognized by her and everyone else. Who should be sent? Who should receive the good of being left behind? It could plausibly be thought that the right thing is simply to send the talented person. But it is also very plausible that doing so would be unfair to her, and that fairness requires a lottery to be held amongst all the candidates. These two views are not incompatible. It may be that fairness requires a lottery, so that it would be unfair not to hold one, but that in this case fairness is outweighed by expediency, so that on balance it is right to send the talented candidate without a lottery. This depends on the circumstances. If it is vital that there should be no slip in the execution of the mission, the unfairness will be tolerable. But if a less than perfect performance is acceptable, more importance can be given to fairness. In some circumstances, fairness will win, and a lottery should be held.

III

Those, then, are the facts. How can they be explained? In this section and the next, I shall present my theory of fairness. I mean

it to apply to the distribution of any sort of good, whether indivisible or not. In Section V, I shall come back to indivisible goods and lotteries.

When a good is to be distributed, for each candidate there are reasons why she should have some of it. These reasons together determine what ought to be done: how the good should be distributed. But how, exactly, do the reasons combine together to determine what ought to be done?² As I shall put it: how do reasons *work*? There are various views about this.

One is *teleology*. Teleology claims that the good ought to be distributed in whatever way maximizes overall benefit.² So the only sort of reasons it recognizes for a particular candidate to get the good is a benefit that would result. Imagine the good being distributed one unit at a time. Each unit should go to the candidate whose receiving it would produce the most benefit; this will normally ensure that overall benefit is maximized when all the units are eventually distributed. At each stage, the reason for giving a particular unit to one candidate is the benefit that would result; the reason for giving it to the next candidate is the benefit that would result from that; and so on. All these reasons should be weighed against each other, and the unit allocated to the candidate for whom the reason is strongest. So we can say that reasons are combined together by *weighing up*. This is how reasons work in teleology. Weighing up goes along with maximizing.

Other views disagree with teleology. One, for instance, claims that some reasons are *side constraints*. A side constraint determines directly what ought to be done; it is not subject to being weighed against other reasons. *Rights* are often thought to be side constraints. Suppose that, amongst the candidates for a good, one has a right to some part of it. Suppose it is her income, for instance, which she has earned. Then side-constraint theory says simply that she should have it; no question of weighing arises. The theory may acknowledge the existence of teleological reasons too, which work by weighing up. It may allow that weighing up is appropriate amongst other candidates, but not for a candidate who has a right.

² There are non-maximizing versions of teleology; see Michael Slote, 'Satisficing consequentialism', *Aristotelian Society Supplementary Volume*, 58 (1984), pp. 139–63. But for simplicity I shall ignore them here.

I am going to describe a third type of reason, which works in a third way. To introduce it, I shall first draw a distinction of a different sort amongst the reasons why a candidate should get the good: some of these reasons are duties *owed to the candidate herself*, and others are not. I shall call the former *claims* that the candidate has to the good.

The distinction between claims and other reasons is easy to grasp intuitively. Take the dangerous mission, for example. One candidate is more talented than the others. This is a reason for allotting to the others the good of staying behind. But the other candidates' lack of talent gives them no *claim* to this good. It may be right to leave them behind, but it is not owed *them* to do so. Whatever claim they have to this good, the talented candidate has it also.

The distinction can appear even within teleology—indeed within utilitarianism. All utilitarians think that if a person would benefit from having some particular good, that is a reason why she should have it. But some utilitarians think this reason is a duty owed to the person—a claim—and others think it is not. William Godwin, for one, thought it was a claim. 'Every man', he said, 'has a right to that, the exclusive possession of which being awarded to him, a greater sum of benefit or pleasure will result than could have arisen from its being otherwise appropriated'.³

The difference is nicely brought out by the attitude of utilitarians to changes in the world's population. Henry Sidgwick⁴ believed an action was right if it maximized the total of good enjoyed by people in the world. So he believed one should promote growth in population if the extra people brought into existence will have good lives, and no harm will be done to people already living. But this is clearly not a duty owed to the people who will be brought into existence. One cannot owe anyone a duty to bring her into existence, because failing in such a duty would not be failing anyone. Sidgwick, then, evidently thought that the duty to benefit people is not owed to those people themselves. On the other hand, a utilitarian view promulgated by Jan Narveson⁵ is that one should promote the

³ *Political Justice*, Penguin Edition, 1976, p. 703.

⁴ *The Methods of Ethics*, Macmillan, 1907, pp. 414–6.

⁵ 'Utilitarianism and new generations', *Mind*, 76 (1967), pp. 62–72.

good of existing people or people who will exist, but there is no reason to increase the total of good in the world simply for its own sake. So the fact that the extra people will enjoy good lives is no reason to increase the world's population. Narveson is evidently motivated by the thought that, whatever duty there is to promote a person's good, it must be owed to the person herself. Consequently, there can be no duty to bring a person into existence.

It is clear, then, that there is a distinction between claims and other reasons. It is not so clear which particular reasons are claims and which are not. Even utilitarians, I have been saying, disagree about this. And if we recognize nonutilitarian reasons, there is further scope for disagreement. If we accept *desert* as a reason why a person should have a good, it is perhaps an uncontroversial further step to take it as a claim. But *need* is more controversial. If a person could benefit from a good, that is no doubt a reason why she should have it, but, despite Godwin, we may be reluctant to accept it is a claim. If, however, the person needs the good, perhaps we should accept that. Perhaps, for instance, a person who needs a kidney has a claim on it. But this is controversial.

In this paper, I am not going to engage in controversy over which reasons are claims and which are not. I shall take it for granted that some are: that some reasons why a person should have a good are duties owed to the person. And I shall concentrate on asking how these reasons, whichever they are, *work*. How do claims combine with each other and with other reasons, in determining what should be done?

IV

Some teleologists, as I say, recognize the existence of claims. But they suppose claims work by weighing up, just like other reasons. They think that the right thing to do, and the right way to distribute a good, is determined by the balance of reasons, whether claims or not. They throw claims and other reasons all together on to the same scales, in the same maximizing calculation.

But the fact that conflicting claims are duties owed to different people gives rise to an alternative intuition. Simply weighing claims against each other may not seem enough. Weighing up is the treatment we would naturally give to conflicting duties

owed to a single person. Applying it between different people may not seem to be giving proper recognition to the people's separateness.⁶

In particular, weighing up claims does not seem to give proper attention to *fairness*. Take the example of the dangerous mission again. All reasons are evenly balanced, apart from the special reason for sending the talented candidate: she will perform the mission better. So weighing up reasons must conclude in favour of sending this candidate. But that seems unfair to her. It might be the right thing to do under pressure of expediency, but nevertheless it seems unfair. The talented candidate has a claim to the good of being left behind, and her claim is as strong as anyone else's. Yet when it is weighed against other people's claims, and the further reason that she will perform the mission better, her claim is overridden. Weighing up seems to override claims, rather than respect them.

It is fairness that matters here because the particular business of fairness is to mediate between the conflicting claims of different people. But I need to qualify this remark slightly. Certainly, fairness is *only* concerned with claims, and not with other reasons. Suppose there is some reason why a person should have a good, but she has no claim to it. Then if she does not get the good, that may be wrong, but she suffers no unfairness. It cannot be unfair to deny her a good she had no claim to in the first place. On the other hand, it is possible that some claims are outside the domain of fairness, and work in different ways from the one I shall be describing. I shall say more about this possibility later, and for the time being I shall ignore it. I shall assume that all claims are mediated by fairness.

Weighing up claims is not enough, then, because it does not give proper attention to fairness. It would not even be enough to give claims extra *heavy* weight in the course of weighing up. The example of the dangerous mission shows this too. However much weight is given to claims, each person's claim to the good of staying behind is still the same. So the claims will all balance, and the talented person will still be sent, because of the extra reason. But this is unfair to her.

⁶ The *locus classicus* for this view is John Rawls, *A Theory of Justice*, Oxford University Press, 1972, pp. 22-27.

What, then, *does* fairness require? It requires, I suggest, that *claims should be satisfied in proportion to their strength*. I do not mean ‘proportion’ to be taken too precisely. But I do mean that equal claims require equal satisfaction, that stronger claims require more satisfaction than weaker ones, and also—very importantly—that weaker claims require some satisfaction. Weaker claims must not simply be overridden by stronger ones.

This suggestion merely extends and tightens up a principle that is often taken for granted: that people identically situated should be treated identically. Economists call this the principle of ‘horizontal equity’.⁷ It is, like my generalization of it, inconsistent with teleological maximizing. To see why, imagine two people have equal claims to some good, but that, if the good is divided between them, less benefit will be produced in total than if it is all given to one. Then maximizing implies it should all go to one, but horizontal equity says it should be divided.

The heart of my suggestion is that fairness is concerned only with how well each person’s claim is satisfied *compared with* how well other people’s are satisfied. It is concerned only with relative satisfaction, not absolute satisfaction. Take a case where all the candidates for a good have claims of equal strength. Then fairness requires equality in satisfaction. So if all the candidates get the same quantity of the good, then fairness has been perfectly achieved, even if they get very little, or indeed none at all.

To be sure, all is not well if they get none at all. For each claimant there is at least one reason why she should have some of the good: the reason that constitutes her claim. Claims should be satisfied, therefore. But it is not *unfair* if they are not, provided everyone is treated proportionally.

Everyone’s claim to a good should, *prima facie*, be satisfied. Indeed, if there is any reason, whether a claim or not, for a person to have some of a good, she should have some. Call this the ‘satisfaction requirement’. Normally, this requirement cannot be fully met for everyone. What does it require then? I suggest it requires maximizing of satisfaction. This implies that, to meet this requirement, claims will have to be weighed against

⁷ See, for instance, Anthony B. Atkinson and Joseph E. Stiglitz, *Lectures on Public Economics*, McGraw-Hill, 1980, pp. 353-5.

each other and against other reasons; I think weighing up and maximizing are appropriate for the satisfaction requirement. But then *also* fairness requires that claims should be satisfied in proportion to their strength. Claims, therefore, give rise to two separate requirements: they should be satisfied, and they should be satisfied proportionally.

It will normally be impossible to fulfil both requirements completely. Consequently, the two will themselves have to be combined together in some way, to determine what should be done, all things considered. Here again, I suggest that weighing up is appropriate: the demands of fairness should be weighed against the demands of overall satisfaction. In some circumstances, no doubt, it will be very important to be fair, and in others fairness may be outweighed by expediency.

In summary, claims work like this. Together with other reasons, they go to determine the satisfaction requirement by weighing up. And claims together determine the fairness requirement by the proportionality rule. Then the fairness requirement is itself weighed against the satisfaction requirement.

Evidently, claims in my theory do not work as side constraints; they do not necessarily prevail. This may be a limitation of the theory. I defined claims as duties owed to people, and it may be that within this class there are some claims that are genuinely side constraints. If some claims are side constraints, they are not covered by my theory. My theory is limited to the subclass of claims that work in the way I have described. Call these 'fairness-claims'. It might be a convenient piece of terminology to say that fairness is a subdivision of *justice*, and that justice is concerned with all claims, but fairness only with fairness-claims.

Consequently, I cannot pretend to have defined claims independently of the notion of fairness, and then shown how fairness applies to them. The subclass of claims I am talking about is partly identified by the way they work, and this is itself determined by the theory of fairness. Nevertheless, I believe the subclass of fairness-claims picked out this way is an important one. It may even include all claims. And for brevity I shall continue to use the term 'claim' for fairness-claims only.

The merit of the theory is that it shows a way claims can work, without simply being weighed up in the manner of teleology,

and also without being treated as side constraints. Robert Nozick argues for side-constraint theory largely on the grounds that teleology is mistaken. He concedes that these grounds are inadequate if there is a third alternative.⁸ My theory of fairness offers one.

It shows how a claim can stop short of a *right*, considered as a side constraint. This fills a significant gap. It seems implausible that anyone has a right to a research grant from, say, the Ford Foundation. But if the Ford Foundation decides to distribute research grants, it should surely deal fairly with the applicants. Someone who was rejected on inadequate grounds would have a just complaint. But how can this be, if she had no right to a grant in the first place? My theory explains how. If her application is good enough, she has a fairness-claim. Consequently, *if* other people are receiving a grant, she should receive one too. In 'Claims of need',⁹ David Wiggins considers just what sort of a claim is generated by need. Again, it seems implausible that a person has a right to whatever she needs. So what can her claim be? I suggest it might be a fairness-claim, which implies that *if* needed resources are being distributed, the person should have a share.

V

Now let us concentrate once more on cases where the good to be distributed is indivisible, and there is not enough to go round.

Take a case, first, where all candidates have equal claims. It would be possible to satisfy their claims equally, as fairness requires, by denying the good to all of them. There may be occasions when it is so important to be fair that this is the right thing to do. But it would totally fail to meet the satisfaction requirement, and normally the demands of fairness will not be enough to outweigh this requirement completely. It will be better to use as much of the good as is available.

In that case, the candidates' claims cannot all be equally satisfied, because some candidates will get the good and others will not. So some unfairness is inevitable. But a sort of partial equality in satisfaction can be achieved. Each person can be

⁸ *Anarchy, State and Utopia*, Basic Books, 1974, note on p. 29.

⁹ In Ted Honderich (ed.), *Morality and Objectivity*, Routledge and Kegan Paul, 1984, pp. 149-202.

given a sort of surrogate satisfaction. By holding a lottery, each can be given an equal *chance* of getting the good. This is not a perfect fairness, but it meets the requirement of fairness to some extent.

It does so, of course, only if giving a person a chance of getting the good counts as a surrogate satisfaction of her claim. This seems plausible to me. After all, if you have a chance of getting the good you may actually get it. It is quite different from merely giving the claim its proper weight against other reasons; that does not satisfy it in any way. Suppose, in the example of the dangerous mission, that the talented candidate was sent because of her talents. She could make the following complaint. She has as strong a claim to staying behind as anybody else. Her claim was weighed against other reasons. But this overrode her claim rather than satisfied it. It was never on the cards that she might actually get the good she has a claim to. But if she was sent because a lottery is held and she lost, she could make no such complaint.

Next, take a case where several people have claims to a good that are roughly, but not exactly, equal. Perhaps, for instance, they all need the good, but not exactly equally. And suppose again that there is not enough to go round them all. Fairness requires satisfaction in proportion to their claims. So if the good goes to the people with the strongest claims, the others will not have been fairly treated; their claims will have been overridden. And if it goes to other people, the unfairness will be worse. So unfairness is once again inevitable. But once again it can, if the circumstances are right, be mitigated by giving everyone a chance of getting the good. Ideally, each person's chance should be in proportion to the strength of her claim: the lottery should be unequally weighted. (At first, it is particularly puzzling how a weighted lottery could be fair.¹⁰ If it is fair for some people to have a greater chance than others, that means they more ought to have the good. So why not let them have it without a lottery? My theory explains why not.) But even a lottery at equal odds may be fairer than giving the good directly to the candidates with the strongest claims. This depends on a complicated

¹⁰ In 'Taming chance', *The Tanner Lectures*, Volume 9, Utah University Press, Jon Elster mentions two examples where weighted lotteries have been used in practice.

judgement. The result of a lottery will generally be that the good goes to candidates who do not have the strongest claims. This is less fair than the result of giving it directly to those who do. The likelihood of this less fair result will have to be weighed against the contribution to fairness of the lottery itself. But it is clear that, if claims are close to equality, holding a lottery will be fairer than not.

A subsidiary point. We have agreed that fairness requires everyone to have an equal chance when their claims are exactly equal. Then it is implausible it should require some people to have no chance at all when their claims fall only a little below equality.

When claims are equal or roughly equal, then, a lottery is *fair*. Whether it is *right* to hold one is then a matter of weighing the fairness it achieves against the likelihood that it will not meet the satisfaction requirement, which in this case requires the best candidates to be selected. The conclusion will depend on how important fairness is in the circumstances. But there will certainly be some circumstances where it is better to hold a lottery than to choose the best candidates deliberately.

A lottery should be held when, first, it is important to be fair and, second, the candidates' claims are equal or roughly equal. These conditions may occur quite often. They do not require an exact balance of all considerations; claims may be equal or roughly equal even when other considerations are not balanced at all. Consider, for instance, life-saving medical treatment such as kidney replacement. It seems plausible that, in these matters of life and death, fairness is particularly important. And it seems plausible that everyone has a claim to life, even if on other grounds some are much better candidates than others. Maybe older candidates have weaker claims than younger, since they have already received a greater share of life. But even so, the candidates' claims may be nearly enough equal to make a lottery appropriate.¹¹ This explains why a lottery may be better than the rule of picking the youngest. If an older person has a claim to the treatment, even if it is a weaker claim than a

¹¹ The arguments of Lewis Kornhauser and Lawrence Sager in 'Just lotteries', *Social Science Information*, are closely parallel to mine in many ways. The main difference is that their arguments permit these authors to recommend a lottery only when claims are exactly equal.

younger person's, it demands proportional satisfaction. A lottery provides at least a surrogate satisfaction: a chance. But the rule of picking the youngest gives no sort of satisfaction at all. It simply overrides the claims of older people. So it is less fair.

That is how my theory of fairness explains the value of lotteries. It satisfactorily accounts for the facts set out in Section II.

VI

I know no alternative theory that explains the value of lotteries as successfully as mine. I cannot review all the alternatives here, but I do need to deal with one that may seem promising at first.¹²

In Section I, I spoke of the 'best' candidates for a good, as though the notion was clear cut. But when the judgement between candidates depends on a comparing reasons of different sorts, it often seems impossible to weigh the reasons against each other in a precise way. How, for instance, when comparing candidates for life-saving treatment, should one's *joie de vivre* be weighed against another's family responsibilities? The impossibility might be in the nature of things: some reasons might simply be incommensurable with each other. Or it might be practical: we might have no practical way of making the comparison accurately, even though in principle the reasons might be commensurable.

This indeterminacy suggests the following defence of lotteries. A group of candidates might not be exactly tied—all equally as good as the others—but even so none might be definitely a better candidate than the others. Or it may be that some members of the group are actually better candidates than others, but we cannot in practice know which. Suppose there is only enough of the good for some of this group. Then, just as a lottery has a natural role as a tie breaker when there is an exact tie, it may be appropriate here for the same reason. Here we have a tie within the limits of comparability. And whereas I said in Section II that an exact tie will be very rare, a tie within the limits of comparability may be common. For life-saving medical treatment, for instance, once the medically unsuitable candidates, and

¹² An argument like this is used by Jonathan Glover in *Causing Death and Saving Lives*, Penguin, 1977, pp. 203–27, and by Nicholas Rescher in 'The allocation of exotic life saving therapy', *Ethics*, 79 (1969), pp. 173–86.

perhaps the very old, have been eliminated, it may be that all the remaining candidates are tied within the limits of comparability. We may not be competent to judge between them.

However, I do not think this argument accounts adequately for the value of lotteries. For one thing, it does not explain their *fairness*. As I said in Section II, the role of a tie breaker is simply as a device for getting the decision made. No question of fairness need arise. The new argument merely extends this role to a wider domain: a device is needed to get the decision made when the weighing up of reasons has gone as far as it can.

Furthermore, if there is no separate reason why a lottery is fair, it is doubtful that it is even going to be the best way of breaking the tie. A fixed rule may well do better, for the reason I gave in Section I: it has some tendency to pick better candidates. Picking the youngest may well do better in the case of life-saving treatment. To be sure, this reason is more questionable in our present context. I said in Section I that, in an arbitrary group of candidates for life saving, the younger ones are more likely to be better candidates than the older. Now, though, we are not dealing with an arbitrary group, but with a group that is tied within the limits of comparability. A candidate's youth is one of the considerations that should already have been taken into account in admitting her to this group. Within the group, therefore, the younger people should generally have fewer other considerations in their favour than the older ones; on balance they should be more likely to be good candidates. However, I doubt that in practice, when a lottery is defended on grounds of incomparability, it will often be for a group of people chosen in this finely balanced way. For instance, the group may consist of all the candidates except the medically unsuitable and the very old. For such a group, picking the youngest would be a better tie breaker than a lottery. In any case, this argument from incomparability provides no reason why a fixed rule should be *worse* than a lottery. It fails the test I mentioned at the end of Section I.

Department of Economics
University of Bristol
40 Berkeley Square
Bristol BS8 1HY.