Selecting People Randomly*

John Broome

I

On the night of April 19, 1841, the American ship *William Brown* hit an iceberg and sank. Two boats got away. One was the longboat containing forty-one people, passengers and crew. It was leaky and so overloaded that no one believed it could survive. After twenty-four hours the weather got worse and the boat was on the point of sinking. The ship's mate, who had been put in charge by the captain, ordered the crew to throw overboard some of the passengers. He was obeyed, but only when he gave the order a second time. Except for a small boy and two men whose wives were with them, all the male passengers were thrown out—fourteen men altogether. None of the crew were thrown out. No lots were cast. The boat was saved by this action and was picked up the next day. A seaman called Holmes, who had helped in throwing the men out, was later convicted of manslaughter. Judge Baldwin, in summing up at his trial, said that passengers ought not to have been sacrificed before sailors and that in any case the victims should have been chosen by lot. "This mode [of selection]," he said, "is resorted to as the fairest mode, and, in some sort, as an appeal to God, for selection of the victim. . . . In no other than this or some like way are those having equal rights put upon an equal footing, and in no other way is it possible to guard against partiality and oppression."[1]

In November 1969 (while the Vietnam War was still in progress) Congress passed a bill allowing a lottery to be used in selecting conscripts for the armed forces. There was very little opposition. In his message to Congress about this bill President Nixon said: "Since more men are classified as available for service each year than are required to fill current

1. United States v. Holmes, 26 Federal Cases 360–369.

or anticipated draft calls, Selective Service Boards must have some way of knowing whom to call first, whom to call second and whom not to call at all. There must be some fair method of determining the sequence of induction for those available for service in the prime age group. In my judgment, a fair system is one which randomizes by lot the order of selection. Each person in the prime age group should have the same chance of appearing at the top of the draft list, at the bottom, or somewhere in the middle." Nixon also said, "Any system which selects only some from a pool of many will inevitably have some elements of inequity."[2]

Chronic hemodialysis is a technique that can save the life of a person whose kidneys have failed. In many countries there are still not enough facilities available to treat everybody who could benefit. Without treatment a patient quickly dies. Different hospitals have adopted different methods of selecting whom to treat. Some have taken into account a candidate's "social worth." For instance, B. H. Scribner in Seattle, who originated the technique, is said to have selected people on these criteria: they should be under forty, free from cardiovascular disease, pillars of the community, and contributors to the community's economics. He also took into account whether they were married and whether they went to church.[3] Later, when the Seattle Artificial Kidney Center was set up, the job of choosing among medically suitable candidates was passed to the anonymous Admissions Advisory Committee of doctors and lay people.[4] At many other hospitals, on the other hand, the procedure has been "first come first served," and at least one has employed a formal lottery.[5] Several commentators have recommended a process that relies on a lottery.[6] Paul Freund, commenting on the Seattle committee, said, "In a matter of choosing for life or death, not involving specific wrong-doing, no one should assume the responsibility of judging comparative worthiness to live on the basis of unfocused criteria of virtue or social usefulness, and . . . either priority in time, or a lottery, or a mechanical selection on the basis of age should be followed. . . . The more nearly total is the estimate to be made of an individual, and the more nearly the consequence determines life and death, the more unfit the judgment becomes for human

2. November 19, 1969, *Congressional Record*, 115, pt. 26:34836.

3. G. E. Schreiner, "Problems of Ethics in Relation to Haemodialysis and Transplantation," in *Ethics in Medical Progress*, ed. G. E. W. Wolstenholme and Maureen O'Connor (London: Churchill Press, 1966), pp. 126–33.

4. J. S. Murray et al., "A Community Hemodialysis Center for the Treatment of Chronic Uremia," *Transactions: American Society for Artificial Internal Organs* 8 (1962): 315–18.

5. Guido Calabresi and Philip Bobbitt, *Tragic Choices* (New York: W. W. Norton & Co., 1978), pp. 178–99; "Scarce Medical Resources," *Columbia Law Review* 69 (1969): 620–92, p. 660; David Sanders and Jesse Dukeminier, "Medical Advance and Legal Lag: Hemodialysis and Kidney Transplantation," *UCLA Law Review* 90 (1967–68): 357–419, p. 372.

6. Al Katz, "Process Design for Selection of Hemodialysis and Organ Transplant Recipients," *Buffalo Law Review* 20 (1972–73): 373–418; N. Rescher, "The Allocation of Exotic Lifesaving Therapy," *Ethics* 79 (1969): 173–86.

reckoning. . . . Randomness as a moral principle deserves serious study. . . . Not even science has taught us how to act as gods on the Day of Judgment."[7]

If a good or bad cannot be distributed equally, it sometimes seems a good idea at least to distribute it randomly. Randomness appears to be a way of bringing some fairness into an inherently unfair situation. This is an attractive thought that deserves careful attention. This article considers what justification can be found for selecting randomly and in what circumstances it applies. These are important practical questions, and because the aim of random selection is to be fair, answering them will also tell us something about the nature of fairness.[8]

One obvious role for random selection is to limit the opportunities for corruption and prejudice in decision making, to "guard against partiality and oppression," as Judge Baldwin put it. But I shall not be concerned with this role. I shall be asking whether, even if there is no question of corruption or prejudice, it is still best to select randomly on some occasions. In Sections III, IV, and V, I shall be presenting what I believe is the best case that can be made out for thinking that it is. In Section VI, I shall argue that this case, even if it is correct, cannot be carried as far as many people have thought. And in Section VII, I shall point out a serious threat to the case itself. Meanwhile, in Section II, I shall mention and set aside an approach to justifying random selection that seems to me inadequate.

II

If we are to produce a significant case for selecting randomly between people, it is plain, first of all, that it cannot depend on there being an exact balance of all considerations. When choosing between people it will never in practice happen that all the considerations in favor of one candidate will exactly balance those in favor of another. So if that were the only time we should select randomly, we should never in practice need to do so. But, more importantly, the value we could attach to random selection, if ever it did become appropriate, would be insignificant. If all considerations balanced exactly, it simply would not matter whom we chose. Or if we did prefer to choose randomly, our preference would be so weak as to be lexicographically dominated by other values; the slightest change in any other consideration would make us give up random selection and deliberately select one of the candidates.

The problem, then, in justifying random selection is to explain how it can be justified even when all considerations are not in balance. One attempt at explanation points to straightforwardly utilitarian benefits of

7. Paul A. Freund, introduction to *Experimentation on Human Subjects,* ed. Paul A. Freund (New York: George Braziller, 1970), pp. xvii–xviii.

8. It also has some important implications for welfare economics, which I have discussed in "Uncertainty and Fairness," *Economic Journal,* vol. 94 (September 1984).

selecting people randomly rather than on merit. These benefits fall roughly into two classes. First, there are matters of cost. As a way of reaching a decision, random selection is effective and cheap. It is cheap in resources, whereas to judge people's merit a lot of effort is needed in collecting information and assessing it.[9] It is also cheap in the emotional burden it puts on the decision makers, whereas to have to judge a person's worth—if the decision is one of life or death—is an awesome responsibility to bear.[10] The other class of benefits are the effects of selecting randomly. It may be that a rejected candidate will be less distressed if the rejection is the result of chance rather than of a judgment about his worth.[11] It may also be that a society that refuses to classify people according to their worth will be a better society to live in.[12] And random selection may command more general assent than other methods and so cause less upset and disturbance.[13]

Jonathan Glover, particularly, lays stress on these utilitarian benefits in justifying random selection.[14] In one way they certainly have the right form. They present advantages of the random process itself which can be set against considerations relating to its outcome. So they do not require the latter to be exactly balanced. Furthermore, these utilitarian points are certainly important. The question of cost, including the emotional cost to the decision makers, might at first sight seem insignificant when people's lives are at stake, but the story of the Admissions Advisory Committee at Seattle suggests it is not. As the medical technology improved, this lay committee was given an increasing share in selecting patients for hemodialysis. "It became progressively apparent, however, that the Admissions Advisory Committee was not able to function effectively under this increased responsibility. The members were slow to take any action and fraught with guilt. In most cases they eventually asked the doctors which of the patients they thought would have the best chance of performing well under treatment. In the summer of 1967 the lay panel was discontinued."[15] Costs, then, can be so high that the system stops working.

However, though important, I do not think these utilitarian benefits can be enough by themselves to justify random selection. The point about cost, to take it first, presupposes that random selection is fair. There are other ways of selecting people that are equally cheap and effective. For instance, the good could be charged for, and the price set so high that only the right number of people are willing or able to pay it. Sometimes

9. "Patient Selection for Artificial and Transplanted Organs," *Harvard Law Review* 82 (1969): 1322–42, pp. 1330–31.

10. Jonathan Glover, *Causing Death and Saving Lives* (Harmondsworth: Penguin Books, 1977), p. 219.

11. Rescher, p. 184.

12. Glover, p. 218.

13. I owe this point to James Griffin.

14. Glover, pp. 218–19.

15. "Scarce Medical Resources," p. 661, quoting an interview with Dr. Tom Sawyer.

it might be possible to select the right number by picking only the whites or only the men. Evidently the case for random selection on a particular occasion must be that it is fair as well as cheap.

The point about effects also works primarily because random selection is fair. If random selection causes less distress to the rejected candidates, if it helps to make a better society to live in, and if it commands general assent, that must be chiefly because it is fair. For instance, it must be less distressing to be rejected by a fair process than an unfair one. It is probably true that not all the beneficial effects result from fairness. Think of a case where selection on merit is undoubtedly fairer than random selection: choosing among candidates for a job. Even in this case random selection, though less fair, might have some beneficial effects. If the job is given to someone at random and you are rejected, you have, perhaps unfairly, suffered a loss. But at least your ability has not been questioned. A meritocratic society, where all jobs are invariably awarded on merit, might well be a bad society to live in. People at the bottom would find it hard to maintain their self-respect.[16] All the same, if random selection has beneficial effects even when it is not fair, it undoubtedly has greater ones when it is fair. And when it is not fair it is hard to believe that its beneficial effects could be a sufficient reason for actually adopting it. We should surely only be in favor of selecting randomly when that is a fair thing to do.

So, if we are to justify random selection on a particular occasion, our first task must be to show it is fair on that occasion. That is certainly not something we can take for granted since there are many occasions when it would plainly not be fair to select randomly. Choosing who is to get a job is one; another is choosing whom to punish for some crime. We shall have to sort out occasions when random selection is fair from those when it is not and explain what makes it fair when it is. And if our support for random selection is to go beyond the trivial, our account will have to explain how it can be fair to choose randomly between candidates even when the considerations in favor of the different candidates are not exactly balanced.

That it would be fair to select randomly on some occasion is a necessary condition for it to be right to do so, but not a sufficient one. When random selection is fair, that is by itself a reason for using it, and its fairness may bring in train as supporting reasons the utilitarian benefits I have described. But there may also be countervailing reasons, and we shall need some understanding of this conflict between fairness and other considerations.

Sections III, IV, and V present an account of the fairness of random selection that, I think, meets these requirements.

16. This, at least, was Michael Young's thesis in *The Rise of the Meritocracy* (London: Thames & Hudson, 1958).

III

When selecting people to receive some good, considerations of various sorts may need to be taken into account. In deciding whose life is to be saved by hemodialysis we might need to consider how much the candidates enjoy life, how much they contribute to the community and their families, what gratitude we owe them for past services, how often they go to church, and so on. Among the considerations will be the duties we owe to the candidates individually to treat them fairly. I shall lump all the other considerations together under the heading of "the general good" in order to contrast them with these particular duties of fairness. "General good" may often not be the most appropriate title, and I use it only for convenience.

To illustrate the distinction I am drawing, take this example. Suppose one of the candidates for hemodialysis has, as well as his kidney disease, a second handicap that prevents him from contributing much to the community and makes his life dull and joyless, though still well worth living. Considerations of the general good will incline us to reserve the treatment for other candidates than this one. Under "the general good" here I am including the candidates' own welfare. In a decision like this there is undoubtedly a duty to promote welfare, and this may be a duty owed to the candidates individually, but it is not the duty of fairness, and it is therefore included in the general good. In this case welfare is best promoted by treating, not the handicapped person, but someone who has a better life. Now, one might reasonably accept all this but at the same time think that fairness points in a different direction: that it is unfair to deprive a person of his chance of living just because his life is not a very good one. So far as fairness is concerned, he has as good a claim to treatment as anybody else.

When a good is to be distributed, a duty is owed to each candidate to treat him fairly. Fairness gives a candidate a certain sort of claim to the good, which must be taken into account. We need to be clear about the nature of such a fairness claim. To say a candidate has a fairness claim does not imply he has any general right to the good. If I say, for instance, that fairness gives a handicapped person a claim to treatment by hemodialysis, I do not imply that he or anyone else has a right to expensive treatment paid for by taxpayers. I mean only that, if hemodialysis is available for some people, this person has a claim to be among them. Fairness is only a relative matter, a matter of how one candidate is treated relative to others. A fairness claim exists only in a particular context where a good is to be distributed, and its only significance is in relation to the fairness claims of other candidates. Suppose, for instance, that all the candidates for a good happen to have equal claims. Then so long as they all get the same amount of the good, however large or small the amount may be, they have been treated with perfect fairness and their claims have been perfectly satisfied.

When a good is to be distributed, fairness gives each candidate some claim to the good. In this article I shall not discuss what determines the strength of each candidate's claim. In different circumstances claims will have different sources. Sometimes they may arise from the candidates' needs, from their general rights such as property rights or a right to life, from a debt of gratitude, or from something else. Sometimes claims may be the resultant of several influences. I shall take no view about any of this, though in examples I shall make various assumptions for the sake of illustration. I shall simply take claims as given and ask what follows.

But I do need to make one assumption about fairness and what determines its claims. To see the point of it consider this example. Suppose we are choosing between candidates for a job. Suppose there are no grounds for positive discrimination or anything of the sort, so that the general good requires us to choose the candidate who will do the work best. If we choose the best candidate on these grounds, it might reasonably be argued that we cannot have acted unfairly. This argument comes in two versions. One says that no particular duty of fairness is owed to the candidates at all; they have no claims on us. The other says that we do owe them a duty to be fair, but this duty simply requires us to select the best on grounds of the general good. A way to distinguish these versions is to ask what wrong we do if we do not select the best candidate. We have certainly failed to do the best for the general good, which is one wrong. But have we also done a particular injury to the best candidate, the injury of treating him unfairly?

This example suggests that sometimes when a good is to be distributed it may be that no considerations of fairness arise, or else if they do arise, they may be entirely determined by the general good: fairness may require that the candidates selected should be those whose selection would best promote the general good. The assumption I need to make is that not all cases are like this: sometimes considerations of fairness do arise and are not entirely determined by the general good. (This assumption is amplified in Section V.) Intuition suggests this is so. The case of the handicapped candidate for hemodialysis is one example. Another is this. In 1943 penicillin was still scarce. The U.S. Army in Africa did not have enough for every soldier who needed it. It had to choose to give priority either to the battle wounded or to the victims of venereal disease. The general good favored the latter group because with the aid of penicillin they could be made fit to fight more quickly. It seems one might reasonably accept this but at the same time think that giving priority to the victims of venereal disease would be unfair to the wounded.

A utilitarian, however, might disagree with these intuitions. He might take either of two views. He might think, on the one hand, that we never owe a particular duty to individuals to treat them fairly; our only duty is to promote the general good. Or he might think, on the other hand, that we do have duties of fairness but that they are always fully determined

by the general good.[17] Every case, then, he will liken to the example of choosing who is to get a job, in one of its two versions. Against this utilitarian opinion I am going to offer no arguments. If it is correct, then the case I am about to present in favor of selecting randomly on some occasions will not work. And I believe no alternative case will be available. We require, as I explained in Section II, an account of how random selection can be fair, and fair even when all considerations are not exactly balanced. But according to the utilitarian view I have described there seems to be no way that random selection could be fair except possibly when considerations of the general good, and hence all considerations, are exactly balanced. I believe, then, that if one accepts the utilitarian view one must deny the value of random selection except in trivial cases (or as a device to guard against partiality and oppression). This may indeed be right. But my purpose is to make out the best case I can for random selection and to see what assumptions it rests on. One assumption, then, is that this utilitarian view of fairness is incorrect.

Granted this, fairness and the general good may conflict. When they do, the weight that should be given to each will depend on the occasion. The judicial system selects people to put in prison. It pays some attention to the general good—exemplary sentencing is one instance[18]—but over-whelming importance is attached to fairness: people should have their freedom if fairness requires it, for instance, if they are innocent. Exigencies of war, on the other hand, can give great weight to the general good. The U.S. Army did actually give priority in the use of penicillin to the victims of venereal disease, and this could well have been the right decision.

IV

It may happen on some occasion that fairness gives several people an equal claim to a good, but there is not enough to go round them all. Fairness requires that all these people be treated equally. But they cannot be treated equally because some of them will get the good and others will not. So some unfairness is inevitable (as President Nixon said in his message to Congress about conscription). However (as Nixon also said), we can go some way toward treating the candidates equally: we can give them all an equal chance of getting the good by choosing between them randomly. This is a less than complete equality, but it is equality of a sort, so it satisfies the requirement of fairness to some degree. It also seems to be the best we can do to satisfy it, or at least I am going to

---

17. I have to thank James Griffin for impressing on me this version of the utilitarian view. William Godwin, *Enquiry concerning Political Justice* (Harmondsworth: Penguin Books, 1976), p. 703: "Every man 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."

18. I owe this example to Sanford Kadish.

assume so until Section VII. It seems that only by random selection, as Judge Baldwin said at Holmes's trial, "are those having equal rights put upon an equal footing."

Although the candidates have equal fairness claims, it does not follow that the general good is equally well satisfied whoever gets the good. I assumed in Section III that fairness and the general good do not necessarily go together. If they do not, the general good will give us a reason to favor some candidates over others. The general good would be best promoted by deliberately selecting those candidates, but this would override the fairness claims of the others. Selecting randomly, on the other hand, satisfies the requirements of fairness to some extent but is likely to promote the general good less well. Deciding which on balance is the right way to select will be a matter of weighing up the fairness of one method against the advantage to the general good of the other. The decision will depend on the relative importance in the particular case of fairness and the general good. But if fairness has any importance at all, the fairness of random selection will be worth some sacrifice of the general good whenever people's claims are equal.

This argument for random selection works only because, as I mentioned in Section III, fairness is relative. A fairness claim is relative in the sense that its satisfaction depends on what happens to the owner of the claim relative to other people, not just on what happens to him. When fairness gives each candidate an equal claim to the good, that particularly means that in the good's distribution each should be treated equally with the others. It turns out that there is a way—random selection—of treating them equally to a degree, even though they do not all equally get the good. The fairness achieved by random selection is less than perfect, not because some people do not get the good, but because the treatment is not perfectly equal.

Imagine for a moment that all claims were absolute in the sense that the satisfaction of a person's claim depended only on what happened to the person, without reference to other people. Absolute claims like this could not explain the value of random selection. If all the candidates have equal absolute claims to a good and they cannot all have the good, an equal number of claims will be satisfied whoever gets it. Selecting randomly cannot increase this number. We might try to justify random selection by saying there is a second way to satisfy a person's claim to a good besides his actually getting it: giving him a chance of getting it. This, of course, would only be a partial satisfaction. By selecting randomly we could then achieve two sorts of satisfaction: first everybody has a chance and then some people have the actual good. But this theory (besides being, I think, very implausible) could not explain why a lottery at equal odds is better than a biased one. Improving one person's chances at the expense of someone else's will increase the satisfaction of the former's claim by as much as it decreases the satisfaction of the latter's. (I leave aside, as even more implausible, the possibility that the satisfaction

of a person's claim might be a concave function of his chances.) The trouble is that absolute claims do not have built into them the particular requirement for equality, which is the requirement that random selection particularly satisfies.

The point of my argument for selecting randomly is to show how it might be the right thing to do even when all considerations are not exactly balanced between the candidates. The argument requires only that considerations of fairness are balanced. Since I have assumed that fairness can be independent of the general good, the general good need not be balanced at all. Nor does the argument require that only fairness counts; the general good may be important, but its importance may be outweighed by fairness.

I said in Section II that in practice one would never find a case where all considerations were exactly in balance. But it is not so implausible that considerations of fairness alone might sometimes be in balance. In some situations people's fairness claims may be determined, not by weighing up various factors, but by some principle that simply makes them equal. For instance, the principles of game playing presumably determine that two chess players before a game have an equal claim to play white. Since fairness in chess dominates other considerations, this explains why they must hold a random draw. (Of course a rule of chess requires them to, but this explains the motivation of the rule.) Judge Baldwin evidently believed that each person in a sinking lifeboat (or at least each passenger) has an equal claim to be saved; I imagine he had in mind some principle such as that everyone has an equal right to life. The Judge also evidently believed that in a sinking lifeboat fairness outweighs other considerations such as the passengers' age or marital status. Consequently he favored a lottery. It may also be that every healthy young man has an equal claim to escape conscription and that may justify a random draft.

But these are questions about what determines fairness claims and are beyond the scope of this paper. I have argued that *if* people's claims are equal it will be right to select randomly, provided the occasion is one where fairness is important enough to outweigh the damage done to the general good.

I need to mention one complication before leaving this section. I have been implicitly assuming that the good to be distributed is indivisible. Sometimes we are faced with an indivisibility in the nature of things; the good, for instance, of staying in a lifeboat cannot be divided. But sometimes the problem of selection arises even with a good that could be divided. The good, for instance, of avoiding military service could be divided by conscripting everyone for a short time rather than a few people for a long time. If people's claims to a good are equal, then fairness requires that the good should be divided between them if it can be. That is the only way to treat them with perfect equality. But this may be very damaging to the general good; short-term soldiers are presumably very inefficient. The damage may outweigh the value of fairness, so that it is best not to

make an equal division. Furthermore, the loss in fairness can be mitigated by distributing the undivided good randomly. This achieves some fairness, though, as I have said, a fairness that is less than perfect. Random selection, then, can help to reduce the conflict between fairness and the general good, making it possible to increase one without too much damage to the other.[19]

V

The argument of Section IV required people's claims to be exactly equal. Exact equality is a stringent requirement, and I think it may be relaxed a little.

Consider this example. Suppose some good, say food for a week, is to be given to either of two people but for some reason it cannot be divided. Suppose the claims of fairness are in this case determined by the people's need; perhaps the food has been sent by a charity for distribution to needy people. Suppose to begin with that the two people are equally in need. Then, if what I said in Section IV is correct, and if fairness in this case outweighs considerations of the general good (such considerations, say, as how easy it is to get the food to the people), the choice between the people should be made randomly. Now suppose that one person is slightly more in need than the other. He has the greater claim to the food, so should he be given it without more ado? If he is, will the other person have cause to complain of unfairness? I think he will. When the people were equally needy, it was fair for them each to have a 50 percent chance of getting the food. Now that one person is slightly—perhaps minutely—less needy, it is implausible that it should be fair for him to have no chance at all. Perhaps the fairest arrangement might be to have a lottery with slightly weighted odds, but let us assume that cannot be done. Provided the difference in need is small, this fairest arrangement would be better approximated by a lottery with equal odds than by no lottery at all.

The position seems to be this. Fairness gives each person a claim to the food, one person a slightly stronger claim than the other. What fairness seems to require is that the people's treatment should be in proportion to their claims. Thus, for instance, if the food could be divided, they should each get some but the needier person more. Because it cannot be divided, perfect proportionality of treatment cannot be achieved. The next fairest thing would be to give each person a chance proportional to his claim. To give them both an equal chance is a little less fair than that. An equal lottery is certainly unfair to the needier person because he is entitled to better treatment, not equal treatment. But to give the food directly to the needier person would be more unfair to the less needy because he is entitled to treatment only slightly less good than the needier person's.

19. I owe this point to Bob Goodin.

This account seems plausible for the particular example. But there also seem to be cases that work differently. The case of selecting candidates for a job is one. I assumed before that, insofar as the question of fairness arises at all, what fairness requires is simply that the job should go to the candidate who is best on grounds of the general good. If it does, the rejected candidates have no right to complain of unfairness. Even if another candidate is only a little less good than the best, it would not be treating him unfairly to give the job directly to the best, so there is no case at all for choosing between them randomly. The situation seems not to permit a slight inequality of fairness claims. Two candidates might, perhaps, have exactly equal claims, but if one falls even slightly below the other, he has no claim at all. The best claim annihilates the others.

In Section III, I described what I called a utilitarian view of fairness: that fairness is always determined by the general good. I said that if we are to make out a significant case for random selection we shall have to assume this is wrong. The view I had in mind—now I can make it more precise—is that fairness is determined by the general good in the manner of the job example: when a good is distributed, so long as it goes to the candidates who are best on grounds of the general good, no unfairness is done to anyone. I take it that this would indeed be the view of a utilitarian, insofar as he recognized the requirement of fairness at all. The assumption was not meant to rule out the possibility that fairness is determined always by the general good but sometimes in the manner in which need determines claims in the food example. If this were true, then, as the food example shows, it would sometimes be right to choose randomly between candidates when considerations of the general good were not exactly, but only roughly, balanced between them.

If my account of the food example is correct, it shows that Section IV's argument for random selection can, sometimes at least, be extended from cases where fairness gives the candidates exactly equal claims to cases where it gives them roughly equal claims. This allows random selection a much wider domain. It must be much commoner for claims to be roughly equal than for them to be exactly equal. For one thing, when claims are determined by considerations that are matters of degree, they will never in practice be exactly equal. For another thing, many sources of claims may be intrinsically vague, so that it may be impossible even in principle to say that claims are exactly equal. Need is perhaps an example. Perhaps it could never be true that two people are exactly equally in need, but it may be quite easy to show that their need is roughly equal. One useful strategy for showing that people's claims are roughly equal is something like this. First you show that characteristics in which candidates differ a lot (such as their sex, perhaps) are not relevant to their claims. Then you show that in characteristics that might be relevant they do not differ much. In this way you can avoid deciding precisely what determines claims. This strategy seems to be the one that Nicholas Rescher adopts in his proposal for a method of choosing patients to

receive exotic medical treatments such as hemodialysis.[20] He suggests picking out a roughly homogeneous group of candidates by nonrandom means and then choosing randomly between them. Since the group is roughly homogeneous, their fairness claims will be roughly equal, so if what I have said is right, random selection may well be appropriate.

VI

I have argued that on some occasions it is right to choose between people randomly. The argument had nothing to do with ignorance or human weakness. If we could appeal to God to make the choice, he too should make it randomly; to make a deliberate choice on grounds of the general good or the strength of claims would on these occasions be wrong.

    According to the argument, to justify random selection on a particular occasion one needs to show that people's fairness claims to the good being distributed are equal or roughly equal. But some authors have tried to justify random selection in a different way. They use a negative argument. They point out the absence of a way we can be confident about for discriminating between the candidates. Consequently, they say, we should treat them all equally, as having equal claims to the good being offered. So we should select randomly. This argument does not necessarily say that claims are equal or even roughly equal, only that we should act as though they are. If we could appeal to God, we might have no need for randomness because God would make the actually best choice. Hence, when we choose on our own responsibility, we can seem to be "playing God."

    An extension of this argument, mentioned by Judge Baldwin, makes out that an apparently random selection such as a lottery may actually be a means by which we do appeal to God. If this were true, it would indeed be an excellent reason for holding a lottery when the choice is difficult. But I shall leave this argument aside because I assume it is no longer credible and also because it is not an argument for random selection. A lottery where God chooses the result is not random.

    The negative argument comes in various versions. One, which Jonathan Glover often uses, argues from human fallibility.[21] Often choosing between candidates is very difficult, as it is when the candidates require hemodialysis. There are many possible criteria, and it is hard to know how much each should count and how well each candidate matches up to them. Some criteria that some people have sincerely thought important (such as going to church) seem irrelevant to others. There is a risk of being deluded by our prejudices. There may not even be any single correct way of weighing up the criteria. And so on. We cannot rely on our judgment, and it is safest to act as though everyone's claim is equal. Another version of the negative argument—this seems to be Freund's

20. Rescher, p. 183.
21. Glover, pp. 203–27.

in the remarks of his I quoted in Section I—insists that we have no right to judge other people's worthiness to live, and for that reason we should treat everyone's claims as equal.

I believe that the negative argument is inadequate. If selecting people randomly to receive a good is to be justified, that requires a positive demonstration that people's claims to the good are at least roughly equal. I shall explain why.

An implicit premise of the negative argument is this principle of insufficient reason: if, when we compare people's claims to a good, we have no good reason to think one greater than the others, then we should treat them as equal. It will be helpful to compare the analogous principle in probability theory: if we have no good reason to think one of several possible events more probable than the others, we should take them all to be equally probable. The trouble with the principle in probability theory is that it leads to different and conflicting conclusions, depending on how you divide up the events you want to attach probabilities to.[22] Imagine you are presented with an urn and all you are told about it is that it contains balls. You are to take out one ball, and you want to know the probability that it will turn out to be red or partly red. You might say, "On the one hand it might be red or partly red, or on the other hand it might be neither, and I have no reason to think one of these possibilities more likely than the other." Then the principle of insufficient reason says that the probability of its being red or partly red is a half. Or you might say, "It might be red all over, or it might be partly red, or it might be neither, and I have no reason to think one of these three possibilities more likely than another." Then the principle says that the probability of its being red or partly red is two-thirds. The fact is, of course, that you can say nothing about this probability. Before you can apply the principle of insufficient reason you need to know what is the correct way to divide up the events that could occur. You need to divide them up in such a way that they are all equally likely. But if you know how to do that, then you do not need the principle.

An analogous difficulty afflicts our negative argument about claims. The argument's conclusion is that because of the principle of insufficient reason we should take people's claims to be equal. But their claims to what good? In order to apply the principle we shall first have to decide what good it applies to. There are all sorts of goods we might have equal claims to according to the principle. For instance, we might have equal claims to live a long life, to live through the next four years, to live to the age of sixty-five, to live a happy life, to complete the bringing up of our children, or something else. All these will have different and conflicting implications for how people should be selected for the good that is actually being offered: treatment by hemodialysis, say. If people have equal claims

22. See, for instance, H. E. Kyberg, *Probability and Inductive Logic* (London: Macmillan Co. and Collier-Macmillan, 1970), pp. 31–38.

to a long life, for instance, preference should be given to younger people since that will tend to equalize the lengths of people's lives. If, on the other hand, the principle of insufficient reason is to justify random selection, we shall have to show that the principle applies to the particular good of treatment by hemodialysis—so that we should treat people as having equal claims to this good—rather than to the good of living a long life or something else. But to show this we shall, in effect, have to show that people have equal claims to be treated by hemodialysis. We shall need a positive argument to this effect before the negative argument can get a purchase. And if we have a positive argument, we shall not need a negative one.

Freund suggests that selecting randomly is a way of refusing to judge between people, but it is not. It commits us to one particular judgment between people: that they have (more or less) equal claims to the treatment. Nor is it a way of affirming people's "equal worth as human beings," as several authors have assumed.[23] For one thing, what is in question is fairness claims and not worth, but let us suppose for the sake of argument that worth determines claims in the cases we are considering. Imagine we have two candidates for hemodialysis. One, if treated, will probably live another twenty years or so, the other forty. Suppose we choose between them randomly. Then if worth determines claims, we are implying that twenty years of one person's life is roughly equal in value to forty years of the other's. A year of one person's life is worth roughly twice as much as a year of the other's. This is not affirming their equal worth as human beings. If worth determines claims, it will require a rather careful assessment of these people's relative worth to justify choosing between them randomly.

Selecting randomly is not a way to avoid playing God. If it is playing God to make a responsible choice between people, then it is also playing God to select randomly; once we have hemodialysis machines and not enough of them to go round, we have to play God. When a problem of selection is hard, the proper response is to do our best with it. We must be sensitive to the various criteria and weigh them up as well as we can. We should pay attention to the political methods by which the choice is made; perhaps we should set up a committee to have different points of view represented, or perhaps a single trustworthy person would act with more responsibility. We should do what we can to guard against merely prejudiced decisions. A lot of the worry there has been about deliberate selection results from the performance of the lay committee at Seattle. Apparently they used criteria that seem to many people prejudiced.[24] The moral to draw from this, though, is not that selection should be random but that that was a bad committee. Evidently the choice should not be left to an anonymous committee constituted as that was.

23. Sanders and Dukeminier, p. 375; "Patient Selection for Artificial and Transplanted Organs," p. 1329.
24. Sanders and Dukeminier, pp. 377–78.

Choosing among candidates for lifesaving treatment is not the only hard choice that has to be made. For many other choices there are conflicting criteria we do not know how to weigh. Someone, for instance, has to decide how much money to allocate to the health service or how to allocate money within the health service: some must go toward saving the lives of old people, some to saving babies, some to relieving pain, some to curing mental illness. When a neurosurgeon operates on a tumor, he has to decide whether he should try to make a complete cure at great risk to the patient's life or else simply give him a few more years of life by removing only a part of the tumor. Here is a surgeon, Jason Brice, talking about this decision:

> With something like a tumour, I am more likely to go for bust [a complete cure] if the patient is a young, single man without a great deal of responsibility. But if he is a family man, with three or four young children, it might be better for me to buy him time. With partial removal of the tumour, for example, he might have another five or ten years, by which time his family would be old enough to accept the psychological assault of his loss. In my view, if you are undertaking really dangerous surgery, you must know something about the patient's circumstances in order to assess what risks you are prepared to take. Obviously you can't have a picture in your mind of a sorrowing wife and young children; you would get into an emotional panic. But you must know of their existence. People say that surgeons are a bit cold, a bit autocratic and unsympathetic with patients. I suspect a lot of this is part of the psychological barrier we put up to protect ourselves when making decisions of this kind. Even if we try to hide from it, we have to be emotional, feeling human beings because we are handling other human beings. The moment a man becomes a pure technician, he becomes a poorer surgeon. I sometimes feel that the strain of living with the decisions you have made creates a greater stress than the operation itself. If I go for bust and don't succeed and the patient dies or is seriously disabled, the failure is hard to bear, longer lasting and more difficult to submerge. You keep asking yourself: should I have made a different decision? Could I have done it this way, or that way?[25]

There are many similarities between the decisions Brice has to make and the decisions that have to be made between candidates for a treatment like hemodialysis. His decisions are no less hard and no less like playing God. But for him to resort to randomness in making his decisions would certainly be wrong, and the same is true of other hard decisions. Before one could justifiably choose randomly the alternatives would have to be such that one could justifiably assign equal chances to them, and that rarely happens. But in this respect choosing between people at first seems different. There seems to be a neutral position we can fall back on when

25. Russell Miller, "The Knife Edge," *Sunday Times Magazine* (London) (November 21, 1982), pp. 34, 36.

our reasons for choosing are hard to weigh up: to treat everyone's claims as equal and give everyone an equal chance. But I have explained that this is not a neutral position; it requires justifying as much as any other. So there is no easy way out. We must do our best with these choices as we do with other hard decisions. We shall no doubt make mistakes. But we are likely to do better than a random selection would have done.

VII

I have presented a case for random selection and tested how far it can be pushed. The case rests on assumptions that may easily be doubted. I have already mentioned one: it assumes a separation between fairness and the general good that a utilitarian may easily doubt. Now I want to mention another. It is an assumption I made in Section IV using Judge Baldwin's words: only by random selection "are those having equal rights put upon an equal footing." When people have equal fairness claims to a good, fairness requires them to be treated equally. If they cannot all equally have the good, I assumed that the best that can be done to treat them equally is to give them all an equal chance of getting it. But perhaps there are other ways of treating them equally. These cannot, of course, be as satisfactory ways as giving everyone the good equally. But perhaps they might be as satisfactory as giving everyone an equal chance. In particular, might it not be just as good to give everyone fair and equal consideration in deciding who should be selected? Might this not satisfy the requirement of fairness just as well as random selection does?

    Giving each candidate fair and equal consideration would, of course, lead us to select the candidates who ought to be selected on the balance of considerations. If this can be done with as much fairness as random selection can achieve, it is plainly a better way of making the choice. The only times we should have a use for random selection would be when it is impossible to give everyone fair and equal consideration. Random selection might have a place in a sinking lifeboat. It might have a place in military conscription if we could not trust our political institutions. It would be no more than a device for guarding against partiality and oppression.

    The question is, then, whether giving people fair and equal consideration is really a good enough way of putting them on an equal footing. It seems to me that at least in some cases it is. Here is an example. (It is modeled on one of John Taurek's, though it is by no means the same as his.)[26] A volcanic island is about to explode. People are waiting to be rescued, some at the north end and some at the south. There are more of them at the north. You have the only boat near enough to help, but your boat is too small to pick up even everybody at the south. You are now at sea approaching the island. You have to decide whether to go to the north or the south. Time is so short that when you arrive at either

    26. John M. Taurek, "Should the Numbers Count?" *Philosophy & Public Affairs* 6 (1977): 293–316.

end the people you rescue will have to be selected by lot. Fairness, assume, gives everyone an equal claim to be rescued. Is there any moral merit in going to the north? If you decide to go to the south because, say, it is cheaper on fuel, have you done anything wrong? You could actually arrange to give everyone an equal chance of rescue: you could decide which end to go to by tossing a coin biased in the right proportions towards the north. If my assumption was right that giving everyone an equal chance is the best you can do to satisfy their equal claims, then that is what you should do. If you have no coin to toss, you should at least go the north rather than the south because that is the way to spread the chances most equally. But I cannot persuade myself that this is right. I cannot believe it makes any difference which end you go to, and I think the reason is this. The same number of people will be rescued either way. So the only way you can act wrongly is to act unfairly. If we were worried about prejudice in your choice—if perhaps the people at the two ends of the island belonged to different races—then we should demand you took some formal steps to secure fairness, such as equalizing people's chances. But if there is no doubt of this sort, if in fact we could take it for granted that you would give no unequal consideration to anyone, we should not concern ourselves at all. So here fair and equal consideration seems a good substitute for spreading the chances equally.

However, I think there may also be cases where it is not. Here is an example. Assume that fairness gives every young man an equal claim to escape conscription. Suppose conscripts are selected not randomly but on grounds of the general good; suppose it is the biggest and strongest men that are picked. And suppose that the selection process gives everyone fair and equal consideration; it really does pick the biggest and strongest. Now a conscript might make this complaint of unfairness. He has been given as much consideration as everybody else, but his fairness claim was to exemption, not mere consideration. Not everyone can be exempted, but for him consideration was not even a gesture in the direction of exemption. Because he is big and strong, consideration could never have done him any good at all. It certainly never put him on an equal footing with the small and feeble. It is really a way of overruling his claim, not of satisfying it. Fair and equal consideration, he might say (begging the question), never even gave him a chance. If we are persuaded by this complaint, we need not necessarily think that conscription should be by lot; this may be a case where the general good outweighs the duty to be fair. But we should have to recognize that the nonrandom selection process, though on balance justified, is unfair.

I find both of these examples quite persuasive, but I am not clear what makes the difference between them. I hope they can be reconciled. But in the meantime the example about the island is a serious threat to the idea that selecting randomly is the only way to satisfy the demands of fairness when people's claims are equal. In the end we may be forced to the conclusion that the only merit of random selection is the political one of guarding against partiality and oppression.