From Choice to Chance? – Saving People, Fairness, and Lotteries¹

Forthcoming in The Philosophical Review

Suppose two people, call them A and B, need to be rescued from deadly danger. Nobody else is around to help but you. You can save either A or B, but you cannot save both. (Think of two castaways on the open sea, with enough time for your lifeboat to reach either of them but not both; or think of two patients in urgent need of an organ transplant, with just one donor organ for you to allocate.) All else is equal: no special obligations, no differences in desert, in likelihood of success, etc. Call such cases *conflict cases*. Here is a question: Are you required to give A and B an equal chance to be saved, by holding a lottery (flipping a coin, spinning a wheel, drawing lots)?²

Many authors in ethics, welfare economics and political science answer: "Yes". For further reference, I will say that these authors endorse the Lottery Requirement, i.e. the claim that we are morally obligated to let a fair lottery decide between A and B in conflict cases. In this paper, I defend a skeptical view. I seriously doubt that the Lottery Requirement is true. Where equal claims conflict and all else is equal, a lottery may be a pragmatically effective way to arrive at a decision; and it may be psychologically beneficial *for us* by relieving the burdens of hard choices. Nonetheless, I see no convincing reason to believe that we are

-

¹ For helpful discussions, I thank audiences in Stuttgart, Zürich, Münster, Erlangen, Konstanz, and Bern, especially Monika Betzler, Lars Dänzer, Anette Dufner, Gerhard Ernst, Simon Gaus, Jan Gertken, Weyma Lübbe, Catrin Misselhorn, Martin Peterson, Daniel Ramöller, Wolfgang Spohn, Ralf Stoecker, Alex Voorhoeve. Special thanks for very helpful discussions to Christian Nimtz, Thomas Schmidt and Christian Seidel. I also want to thank an anonymous reviewer for enormously helpful criticisms and suggestions.

² The notion of a lottery that is typically used in the present debate is narrower than the technical notion used, e.g., in expected utility theory. It follows a more intuitive usage according to which not all allocation procedures are lottery procedures. (If we simply identify lotteries with distributions of probabilities over outcomes, as it is customary in expected utility theory, any allocation mechanism will correspond to a lottery.) To capture the narrow, intuitive sense involved here, we can define a lottery as an allocation procedure that is designed to yield a specific distribution of probabilities over outcomes. The precise notion of probability involved here is a tricky matter. However, this is a difficulty to be sorted out by supporters of the view that I am opposing, and I will touch on the issues only where they matter to my argument. When I talk of a fair lottery, I usually mean a device that is designed to yield an equal distribution of probabilities – unless I explicitly note otherwise. (Later on I will discuss some proposals to the effect that lotteries can be fair without affording any kind of equiprobability.) Note that this talk of a fair lottery is descriptive; cf. the typical talk of fair coins. So there is nothing incoherent in maintaining that there is no requirement of fairness to use a lottery while still calling the lottery at issue "fair". ³ E.g., Broome (1984; 1990; 1998), Childress (1970), Daniels (2012), Diamond (1967), Elster (1989), Goodwin (2005), Kamm (1993), Kornhauser & Sager (1988), Saunders (2008), Sher (1980), Stone (2011), Wasserman (1996). Also, see many followers of J. Taurek (1977) (but probably not Taurek himself), referred to in footnotes 26 and 27.

morally required to hold a lottery in conflict cases, as opposed to simply picking one of the parties on more subjective grounds or out of pure whim.

I do not directly argue for the negative claim that there is no moral duty to hold lotteries in conflict cases. I merely try to cast doubt on the defensibility of the positive claim. I argue that our current best attempts at defending the Lottery Requirement fail. This warrants skepticism about the Lottery Requirement; but, of course, there is no proving that no better defenses are possible. Of course, if my argument is convincing it does not follow that we must not hold lotteries. In view of the non-moral advantages they can have, we may reasonably decide to use them. But that does not mean that my argument is without practical import. While the practice of using lotteries may be unproblematic, the practice of treating the Lottery Requirement as a valid moral principle is not. If we treat it as valid, we will, e.g., tend to sanction people for failing to use chance mechanisms in allocation conflicts. My argument will suggest that such practices lack justification.

The structure of this article is straightforward. I present three strategies to defend the Lottery Requirement. Sections 1 to 3 describe and criticize each of these strategies. Section 4 discharges some remaining grounds for resistance to my skeptical conclusions, and it discusses the possibility of defending a weaker version of a normative lottery principle. Finally, section 5 discusses implications for debates about saving the greater number and axiomatic utilitarianism.

Here are the three accounts to be discussed in an overview. According to the first account, holding a lottery in conflict cases is a requirement of distributive fairness. Although it is not a way to ensure that the parties receive equal benefits (since only one party can be saved), a lottery does afford a distribution of equal shares of something that is relevant to the parties' claims. For reasons that will become evident, I call this the *surrogate satisfaction account* (SSA). A second view claims that a lottery is a requirement, not of distributive fairness but of procedural fairness. I thus call it the *procedural account* (PA). A third account argues that in conflict cases, a fair lottery is the unique option that all parties have sufficient reason to consent to, and that this explains why it is morally required. I call the defense of the Lottery Requirement that draws on this idea the *ideal consent account* (ICA).

1. The Surrogate Satisfaction Account

Here is how Kornhauser and Sager (1988) describe the core of the view that I have labeled SSA: "Equiprobable lotteries facilitate equal treatment ex ante. Suppose we have a single good which is indivisible and two claimants of the good who each enjoy equal and robust

moral entitlements to it. An equiprobable lottery has the unique capacity to divide the good probabilistically, making possible the allocation of a ½ chance of receiving the good to each candidate" (ibid., 142).⁴

On the face of it, this is a puzzling idea. How can an indivisible good be divisible after all, and what is it to divide a good probabilistically? After all, giving someone a chance to receive, say, a kidney, is not to give them a kidney "to some degree". Nor is it a way to give them some part of a kidney (whatever good that might do). So what could the defenders of SSA have in mind?

Maybe we have misidentified the relevant claims here? True, it sounds odd to suggest that a chance to some good is the right kind of thing to satisfy a claim to the good itself; but why not say, directly, that what the parties have is a claim *to an equal to chance* to get the good? This will not do. After all, defenders of SSA try to explain why we owe it to both parties to give them a fair chance. It is not an explanation at all to state *that* they have a claim to a fair chance. And in the dialectic of the present paper, appealing to a claim of this kind would beg the question at issue.⁵

So there is a reason why SSA appeals to claims to the good itself, and why Kornhauser and Sager insist that a lottery offers "an equal division of a good that is otherwise indivisible" (ibid., 146). Also, it is independently plausible that what makes a chance to be saved important to the parties is precisely the bearing that it has on the satisfaction of the parties' interest in being saved. So there is something to be said for explaining claims to fair chances to a good in terms of claims to the good itself. SSA offers such an explanation. Still, there is something odd about the proposal that distributing chances to some good somehow amounts to distributing shares of the good itself.

Other presentations seem conscious of the odd ring of the proposal. Here is how Broome puts the idea: "[T]he 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 given a sort of surrogate satisfaction. By holding a lottery, each can be given an equal *chance* of getting the good" (Broome 1990, 119).

⁴ Other defenders of this idea include Broome (1984, 1990, 1998), Eckhoff (1989), and Saunders (2008).

⁵ Strictly speaking, the Lottery Requirement merely holds that there is a duty to hold a lottery. It does not mention *directed duties*, that is, it does not say that we have a duty *to A and B* to hold a lottery (and, correspondingly, that they have a *claim*). Still, I take it that this appeal to directed duties would beg the important question. If we are trying to settle whether we have a duty to hold a lottery, it is not enough to simply say that we owe such a duty *to A* and B.

So the idea is that although you cannot satisfy the claims of both A and B, you can provide both with "a sort of surrogate satisfaction". (Note that the lottery is supposed to achieve a "sort of partial equality of satisfaction" of claims to the good itself. If the idea was that lotteries satisfy an independent claim to an *equal chance* to get the good, they would give full and genuine satisfaction.)

But how can we understand the idea that chances of satisfaction are surrogate satisfaction? I think the best defense of SSA is to start from facts about rational preferences and/or from facts about betterness.⁶ Saunders (2008, 367) writes: "I believe the chance of a benefit is itself a benefit; after all, almost everyone prefers to have a chance rather than none." I agree that people prefer to have a chance rather than none, and to have a higher chance rather than a lesser one. And I add that it is rational for them to do so. I also add that it is better for people to have a chance rather than none, and it is better to have a higher chance rather than a lesser one. So by giving people a fair chance, we give them something (comparatively) *good*, and something they *rationally prefer*. Surely, one may think, a rationally preferable good is a benefit, i.e. the right kind of thing to satisfy the relevant claims?

Actually, this is not clear at all. As I will argue now, the above claims about betterness and rational preferability hold only across a restricted range of cases. In general, the following may be true: As long as we look only at allocation procedures *of a certain type*, we have a reason to prefer those procedures that give us X over those that do not. And within that class, procedures that give us X may also be better for us than those that do not. Crucially, it does not follow from any of this that we have reason to prefer those procedures that give us X over *all* other allocation procedures. Nor does it follow that we are benefitted in any way if one of the procedures that give us X is chosen.

As a counterexample, take certain instrumental goods. Suppose there are several hungry people but only one can of food to allocate. There are two possible procedures. Procedure A is a two-step procedure. In a first round, can openers are distributed. Each potential receiver of the can gets one. In round two, the can is allocated somehow. Procedure B is a one-step procedure. We allocate one can opener along with the can; those who do not get the can do not get an opener. (Suppose we use the same procedure here as in step two of Procedure A.) Clearly, *if* the two-step procedure is chosen, it is better for people, and rational for them to

_

⁶ Deriving claims about goodness from claims of rational preference is, of course, a standard maneuver in economics. (I briefly discuss expected utility theory and axiomatic utilitarianism below.) For an approach that takes relations of betterness (instead of preference) as basic, see Broome (1991).

prefer, that they receive an opener in round one. (No one should say: "I pass." Else how do they open the can they may receive in round two?) But clearly that does *not* entail that people have a reason *to prefer Procedure A over Procedure B*. It would be a mean joke to say that people should prefer Procedure A because that way, at least everybody gets their very own can opener as "surrogate satisfaction".

As long as we consider two-step procedures, potential receivers will have a reason to prefer those procedures that distribute can openers equally over *other* two-step procedures – e.g., those in which not all potential receivers get a can opener in round one. But, as the example shows, it is a misunderstanding to think that it follows that we have a reason to prefer the former two-step procedures over *all* other allocation procedures. Also, it is a mistake to think of the equal distribution of can openers as an equal distribution of substantive benefits, benefits that make people better off compared to those allocation procedures that do not give everyone an opener.

Let me stress that I do not claim that there is a perfect analogy between can openers and chances. Still, similar considerations show that chances given by a lottery do not benefit us in a way that makes it rational for us to prefer, and better for us, that a lottery procedure is chosen instead of certain non-lottery procedures. A lottery that gives you a certain chance can be better and preferable with regard to *some* other procedures simply because by offering you that chance, the lottery has a greater expected value for you. For that to be true, the chances must merely serve as probabilistic weights for the calculation of expected value. They need not be valuable in themselves, and they need not give you any reason to prefer lotteries over *all* other allocation procedures.

To see this, we may first consider an *ex ante* comparison of procedures. Suppose you are B and you need to be saved. You learn that one of the following two procedures will be chosen:

Procedure One: A fair lottery will be held to decide between A and you.

Procedure Two: The rescuer will save the person whose name was written on a piece of paper, without a lottery, earlier.

Suppose you do not know whose name is shown on the piece of paper; but you are *equally* sure that it may be A's name and that it may be yours. That is, your subjective probabilities

with regard to the pessimistic hypotheses and the optimistic hypotheses sum to exactly .5 respectively. Here is a question: Do you have any reason to hope the first procedure is chosen? Would it be rational for you to pay any sum of money to make sure that the first procedure will be chosen?

It seems not. Your expected utility is the same in both cases (at least if your subjective probabilities with regard to Procedure One are fifty-fifty as well, which they rationally ought to be, see Lewis 1986). The lottery makes no difference. In fact, it seems to be *double counting* to regard the chances provided by the lottery mechanism in the first procedure as additional goods that should be added to the expected utility of being saved with a probability of .5. ("First, there is the 50% prospect of *survival*. And then, of course, there is the *chance*.") This clearly suggests that these chances should not be calculated as additional values in an *ex ante* comparison across procedures.

One might reply that the chances that are involved in a fair lottery *just are* the subjective probabilities of the parties that need to be saved, i.e. A and B (you). Therefore, both procedures would yield the relevant kind of equal chance and thus provide surrogate satisfaction. In other words, even Procedure Two would be a fair lottery. But that cannot be true. First, it would mean that it would be impossible for A and B to falsely believe a lottery to be fair. Secondly, it would render the account of surrogate satisfaction hopeless. Your equal subjective probabilities regarding Procedure Two are just a matter of ignorance. And surrogate satisfaction, whatever it may be, cannot plausibly consist merely in being left in the dark about your prospects.

Maybe you, the reader, do not agree that you have no reason to prefer Procedure One, or to be ready to pay a sum of money for it. (If so, this is probably because you have a so-called ambiguity aversion, not because you find that the chance has an extra value over and above the prospect of survival offered by the lottery.⁸) Anyway, I invite you to imagine instead that you are entirely and univocally sure that it is *more* likely than not that your name is written on the paper. In that case, isn't it more rational for you to hope that Procedure Two is chosen? After all, it makes survival more likely by your lights! Still, the defender of SSA must hesitate to answer: "Yes". After all, I did not tell you *how much* more likely you consider it that your name is written on the piece of paper. If that surplus probability is small enough, it should be

_

⁸ What has come to be known as Ellsberg's Paradox (Ellsberg 1961) suggests that many people manifest ambiguity aversion. That is, they prefer gambles in which they have exact knowledge about the odds to gambles which should have an equal expected utility but which involve a greater deal of uncertainty (as opposed to known risk). Of course, this does not show that ambiguity aversion is rational. But since I also do not want to rule out that it may be rational, I consider the following variation of the case.

possible that the additional 'surrogate value' of the chances in the lottery in Procedure One outweighs this. (Once again, note that we cannot say that Procedure Two likewise gives you a chance in the sense that matters for surrogate satisfaction. The only chance that is involved here is a matter of ignorance.) I think this nonsense possibility should convince us that SSA is committed to double counting.

Of course, you might take yourself to have *moral* reasons to prefer Procedure One; maybe you prefer a fair decision procedure to a better prospect of survival. But we cannot at this point appeal to such heroic preferences, and not just because they are likely to be quite rare. One reason is that we are precisely trying to find out *whether there is any ground* for such a moral preference. Another reason is that SSA has a specific form. As I said above, the defense of the Lottery Requirement offered by SSA says that a chance to be saved is a particular kind of good, a good that that answers precisely to *a claim to be saved*. So what we should assume in considering the example is that you look at the two procedures solely with that interest in mind.

These examples discredit the idea that chances are surrogate goods in a relevant sense. The relevant facts about rational preference and betterness are fully explained by the fact that lotteries affect the expected value that a procedure offers; the comparison of Procedures One and Two shows that they are not additional goods in their own right. Especially, these examples show that it is not the case that a fair lottery is better for A and B, and rational for them to prefer, vis-à-vis all other non-lottery allocation procedures.

This verdict finds more support if we turn to an *ex post* perspective. Here, the dominant impression is that when the lottery is over and the time comes to save one party, the previous distribution of chances does nothing to make the parties better off. Here is how C. Hare (2012, 385) puts the point: "If I die then I am not significantly better off for having had a high chance of living. If I live then I am not significantly worse off for having had a low chance of living." If you have doubts, imagine the following. You are B; a coin is tossed; you lose. So you face death. Then, however, the relevant authorities say that the lottery *might* be declared invalid. Instead, a new lottery might be held. In that case you would get a new chance but, for some reason, a lower one. Crucially, if the authorities were to change their mind in this way, this would mean that you never really had a 50% chance. But surely it would be weird for

_

⁹ Is that a weird case of changing the past? No. The fact that some lottery gives someone a certain chance is in part a fact about what happens *after* the chance mechanism has done its work. A lottery gives you a 50 % chance to receive a certain good only if two things hold: the chance of the favorable outcome is 50 % *and* the outcome is afterwards *treated as binding*, i.e., it actually settles whether you get the good. The change of minds affects

you to protest on the ground that this change of minds makes the 'surrogate good' of a 50 % chance you thought you received naught. It seems better for you, and rational for you to prefer, to sacrifice whatever surrogate good you would have received for an arbitrarily low new chance.

These intuitions in our *ex post* assessment confirm that chances that one has received are not substantive benefits. It is not that such chances are goods in their own right that somehow lose all their value by the time the actual saving begins. Their only relevance is as weights in calculating expected value, weights which change to zero or one once the outcome is known.

By way of a general summary: Both the example of the can openers and this last case suggest that certain things have a kind of 'value' that is an artifact of a certain type of allocation procedure and that this value may not carry over to a comparison with certain other procedures.¹⁰

There is a further problem with SSA. So far, I raised doubts about the idea that chances to goods are themselves benefits of a kind that can constitute surrogate satisfaction. The further worry surfaces once we take seriously the fact that the goods in our cases are indivisible in a sense that is not causal but *conceptual*. It is not that we are somehow causally kept from providing people with half a rescue. If you are only 'half saved', you are not saved. So *even if* a 50% chance to a good were sometimes worth half as much as the good – if the good in question is survival, what could it mean for something to be 'half as good' as *that*? If what you need is a kidney, or a lifeboat, then how could anything that is supposedly 'half as good' as that do you any good at all? True, if your preferences are coherent in the sense of von Neumann and Morgenstern (1944), expected utility theory shows that there is a set of functions that cardinally measure preference satisfaction by assigning certain quantities, i.e. "utilities". Chances to be saved will have a certain amount of the "utility" of being saved, which may suggest that it does make sense to say that, e.g. a chance to a kidney is "half as good as" a kidney. But this, I think, just goes to show that it is wrong to reify "utilities", i.e. to turn abstract representations of preferences into real quantities.¹¹

this latter factor. So although the chance of, say, heads was 50%, it turns out that it was not a chance *for you to get the good* after all.

Of course, the instrumental value that the opener has for the one who receives the can is not such an artifact. But the value that can openers have for all potential receivers in round one is.

¹¹ See, for example, Hampton (1994). I think the same worry applies if we follow Broome (1991) and apply expected utility, not to preference relations but to the evaluative relation of betterness that was discussed in the preceding.

I see no relevant sense in which chances of satisfaction are surrogate satisfaction, so that giving people equal chances would constitute a way to treat claims fairly in the distributive sense.

2. The Procedural Account

Even if you cannot give A and B equal shares, you may be morally required to decide between them in the right way. So a second idea is that a lottery is a requirement, not of distributive fairness but of procedural fairness. To defend this approach, one must first claim that there is a certain procedural norm that requires that a fair decision process have certain characteristics. Secondly, one must argue that holding a lottery is a process that has these characteristics. Thirdly, one must show that lotteries are *unique* among the feasible decision procedures in having these characteristics. I will raise various worries about PA in what follows. But my main objection will be that all versions of PA fail at the third stage. On the most plausible views of what a fair decision procedure must be like, other methods besides lotteries have these characteristics, too.

So what is the relevant procedural norm? One seemingly plausible candidate should be discussed first, as it is important to see that it does not apply in conflict cases. One might think that a lottery is required because it shows a kind of *equal concern*, in the sense that it manifests the wish to further the good of both parties equally. This is a mistake. In conflict cases (as defined), the survival of A and the survival of B are not independent goods. The chances you distribute are stochastically dependent: Any chance that A survives is a chance that A *rather than B* survives, and *vice versa*. So should a decent person wish that, say, A survives, and be willing to promote that result? Only as long as she manages to ignore that this means that B will *not* survive, and *vice versa*. Not only does such positive concern for any individual party seem indecent (at least if we stick to our assumption that no special relations to one of the parties are involved). Furthermore, since this concern seems to involve a corresponding neglect of the other party, it seems that having this attitude towards both parties equally involves a clear tension (maybe even an inconsistency).

A much more promising candidate for the procedural norm is of a more negative kind. We may still call it a norm of equal concern, but what is really meant is that a decision should

-

¹² It is easy to confuse conflict cases with the following kind of case: A and B are sick; you have 100 units of a certain drug; each additional unit is known to give patients an additional 1% chance of being cured. Here the chances you distribute will be *stochastically independent*. If you give each 50 units, there is a 25% chance that both are cured. So it is apt to see the cure of A and the cure of B as aims that conflict only insofar as your resources will not allow you to go all the way towards definitely attaining them. Here it *would* make sense to require equal measures of positive concern.

not be distorted by some form of unequal concern. And in fact, there is wide agreement in the literature that the decisive procedural characteristic of a lottery is of this negative kind: It is impartiality. A lottery allows you to make decisions that result in the saving of a person without forcing you to choose this person over the other one. You assign acts (saving A, saving B) to outcomes (heads, tails), run the lottery and perform the selected act. At no point do you have to ask which person to save. In a sense, lotteries establish a kind of veil of ignorance: You make decisions that suffice to select a person without being in a position to tell which one will be selected. This way, lotteries serve to rule out all problematic partial motives. As P. Stone puts it: A lottery is needed "when, and to the extent that, it is important that bad reasons be kept out of the decision. Lotteries can perform this task because they make decisions on the basis of no reasons" (2011, vii). Stone, following Kornhauser and Sager (1988), describes this as the "sanitizing" function of lotteries.

This kind of reasoning is typical for proponents of PA, and I will mostly concentrate on it. One might worry that this line of thought is insufficiently general to lead us to the Lottery Requirement. True, there are many cases in which we have reason to fear that our choice might be influenced by motives that are sexist, racist or otherwise objectionably partial. But surely there may also be cases in which the parties in need share all the characteristics that can reasonably be expected to be relevant. In such cases, particular partial motives simply cannot discriminate among the parties. But the Lottery Requirement is supposed to cover cases like these as well. Surely its defenders will say that people deserve an equal chance to be saved even when the other party happens to be very similar to them. So what is required to make PA work seems to be the claim that there is something objectionable about *any* direct choice of one party over the other.

But maybe this idea is plausible. Consider: If you directly choose A over B, does that not mean that you declare that A should live rather than B? Do you not treat A as somehow more important, or more worthy? According to this line of thought, the problem with a direct choice is not just that it might be grounded in particular problematic motives. Rather, any such choice will be partial in the sense that it implies or conveys a judgment of unequal worth. This may be the general reason why direct choice is impermissible and a lottery mechanism is required.

Let me grant this for now and turn to my main objection to PA. It is that, although a lottery may be sufficient to satisfy requirements of impartiality, it is not necessary. There are other ways to make decisions that are not based on bad reasons. (Ultimately, I will claim that simply picking may do.) Here is a first example. If all that a fair lottery is supposed to do is

keep out motives of the wrong kind, then it should be enough that the decision-maker *believes* that the lottery is fair. Suppose you toss a coin that you and everyone else concerned reasonably believe to be fair. After the lottery is over, some scientists decide to have a look at the coin. Just when you set yourself to save the winning party, they inform you (and the people in need) that the coin was heavily biased and that the losing party did not really get a chance. Crucially, if PA is correct, that should be no reason at all to reconsider. You believed that the coin was fair (and we may suppose that nobody doubts that you believed this), so the veil of ignorance was in full effect. There were no specific bad reasons and no claims of unequal worth.

This is a problem. Intuitively, this was not a fair lottery. Many of those who are sympathetic to the Lottery Requirement will think that the loser has a reason to complain. No matter what you thought, he or she had no chance. Yet the procedure achieves what PA claims should be achieved.

But proponents of PA may have something to say in response to this example. Later on, I will discuss whether they could simply bite the bullet. Maybe they can hold that, although some of our intuitions seem to call for chances of a more objective variety, the notion of a lottery leaves room for interpretation. And maybe there is room to claim that the process described *is* a fair lottery in some adequate sense. (I will call this general type of maneuver the Inclusivist Reply.) For now, however, it is worth pointing out that proponents of PA also have a different reply. They can say that, the example notwithstanding, their account still entails that any *practically applicable* lottery principle requires fair chances in a more objective sense. Though the procedure above satisfies the relevant procedural norm, it is not a procedure we could intentionally apply. This is because we cannot knowingly adopt a false belief. In practice, the only way to make ourselves believe that a process yields equal chances is to convince ourselves as best we can that it *does*. So even according to PA, any procedure we can apply will have to be, by our lights, a fair lottery in the relevant objective sense.

Fair enough. But the problem goes much further than that. There are examples of decision procedures that have the sanitizing characteristic but do not even require the belief that one gives each person an equal chance. Kornhauser and Sager (1988, 140) give an example of such a procedure that uses people's dates of birth as a criterion; I describe a variant of it. If you do not yet know the dates of birth of A and B, you can use some chance mechanism to select a random day of the year (Kornhauser and Sager suggest a 365-sided die) and decide to save the person whose birthday is next counting from the chosen date. Clearly, there is no reason at all to suppose that this gives A and B equal chances. If A's birthday

happens to be a week before B's, then A's chances are vastly better. (As Kornhauser and Sager observe, equal chances for an even number of candidates are in fact impossible unless the candidates were born in a leap year and we therefore use a 366-sided die.)

Nonetheless, the relevant sanitary characteristics (i.e., the veil of ignorance) are again in full effect. Importantly, as Kornhauser and Sager note, the date of birth merely "functions as a *mechanism* for choosing, not a *reason*, since the agency is merely tagging the candidates; there is no thought that a particular candidate deserves to win because of her birthday" (ibid.). So there is nothing in this process that presupposes an unequal worth of persons. Nor is there any way for racist attitudes, sexist attitudes or anything the like to exert an influence on this process.

Defenders of PA cannot in all seriousness object that the agency in this case might have a secret immoral bias towards, or concern for, people who have a date of birth that is close to the number showing on their die's upper side. If we consider *that* a serious possibility, we must do the same for fair lotteries. We must ask whether we can be sure that there is no secret conviction that our coin magically renders chosen candidates more worthy. This is just absurd. Once this is clear, I think the defender of PA should admit that there is not even a need for procedures as complicated as the one suggested by Kornhauser and Sager. Why not just pick the person whose birthday is up next *from today*? Or the person that will make a sound, or blink, next? There is an indefinite number of feasible mechanisms, all of which are, I should think, beyond the suspicion that there is some immoral fetish or bias that implicitly ascribes unequal worth to the candidates.

I think we should even admit that we can simply pick among persons without presupposing that we have reasons to choose one person over another. Generally, in cases in which we see no reason to choose one person over another, we need not end like Buridan's ass. (Otherwise I would never get to ask someone for the time in a crowded street.) When the reasons to pick A or B are equal, or on a par, or incommensurable, we must, and can, simply pick without presupposing that the reasons are *not* equal, or on a par, or incommensurable. Some factors can be triggers for picking rather than reasons for choosing (see Ullmann-Margalit & Morgenbesser 1977). So even picking a person in a conflict case need not involve a judgment of unequal worth.

It may be objected that we can never be sure that we really are 'merely picking.' There may be all kinds of implicit biases that influence our choice. But here it is important to be clear about which biases are really problematic. Suppose I pick A, who wears a very bright shirt. Later, social psychologists find out that on 99 out of 100 occasions, I pick the person

with the brightest shirt. Is that a problem? Is that a mechanism that should have been bypassed by an acceptable decision procedure? I do not see why. Of course, the findings reveal that the parties did not get equal chances. But the defender of the Lottery Requirement cannot object to the mechanism for *that* reason without begging the question. What if A knew of my tendency and decided to wear a bright shirt in order to "nudge" me into picking him?¹³ I do not think this is problematic either. It is not objectionable for parties in need to try to catch the rescuer's attention by shouting, waving, etc. Why would it be wrong if one of them exploited their knowledge of my tendencies?

But what if I myself knew about my bias beforehand, and I consciously decided not to bypass it by using a lottery? Again, I think that this need not be problematic. If I decide not to counteract my tendency to pick the one with the brightest shirt, this does not necessarily mean that I treat the color of people's shirt as a reason. I think it may be just like the case in which I fix an arbitrary criterion (see above), only that I allow my bias to fix the criterion for me. Take this trivial analogy: I need to ask someone for the time in a crowded street, and I decide to pick the one whose shirt I like best. Here, I do not treat my preference as a reason why I should ask that person. I certainly do not take that preference to have any bearing at all on my present concern, which is to find out what time it is. My preference is just a trigger for picking, not a reason for choice. I see no reason to deny that my bias towards bright shirts (or even a preference that is not 'implicit' at all) could play a similar role in the more dramatic cases that we are concerned with.

Of course, things look differently if we take biases towards a particular sex or race, etc. But even here, the explanation of why they are problematic cannot, on pain of circularity, be just that the parties do not get equal chances. Defenders of the Lottery Requirement need a non-question-begging account of which biases are problematic. (One proposal: A long history of racist and sexist discourses, practices and institutions makes it likely that a sensitivity to gender or race, unlike a sensitivity to the color of shirts, is rooted in deep-seated evaluative attitudes that are objectionable in themselves, apart from their effects on chances.) Anyway, there are problematic biases, and I grant that they should be kept from influencing decisions in conflict cases. Note that it is not clear that the appeal to problematic biases will help us in defending the Lottery Requirement in its full generality. As I said, defenders of the Lottery Requirement will maintain that the parties deserve equal chances even if they happen to be similar with regard to features like sex, race, age, ability, etc. But I will bracket this generality

¹³ See, of course, Thaler and Sunstein (2008).

concern. Maybe defenders of the Lottery Requirement may reply that we can never be entirely sure which features are the objects of problematic biases, and that in all realistic instances of conflict cases the parties will differ in *some* features.

So the defenders' claim would have to be that there is a *general suspicion* that we are susceptible to problematic biases, and that this suspicion is serious enough to ground a general principle like the Lottery Requirement. For what it is worth, I do not think that such a sweeping suspicion is warranted. Note that the claim is not just that we must always be wary of hidden problematic biases. With that I agree wholeheartedly. Rather, proponents of PA must claim that the threat of a hidden problematic bias is so virulent, in every single case, that we would always act immorally in placing trust in our ability to merely pick. I am not convinced that the latter is true. But, fortunately, we need not argue about this. For defenders of the Lottery Requirement will find that this general suspicion does them a disservice anyway. As with other skeptical hypotheses, it is hard to keep a lid on the consequences of this general suspicion. If proponents of the Lottery Requirement base their defense on it, they will find themselves driven to a more extreme view.

Let me explain. As defenders of the Lottery Requirement point out, this principle deals only with cases that are *morally symmetrical*, i.e. cases in which both parties have "equal and robust moral entitlements" (Kornhauser and Sager, see above). They certainly do not mean to suggest that it is fair to hold a lottery even where the parties seem unequal with regard to desert, need, etc.

If the morally symmetrical nature of conflict cases matters (which no one seems to doubt), part of what a fair procedure must do is ensure that this symmetry is correctly registered. But people can certainly be partial in judgments about whether the conditions for a lottery that the Lottery Requirement presupposes are satisfied. It is naïve to think that racist and sexist biases will sit and wait until the symmetry of the situation has been established in an impartial way. Rather, they will distort our prior judgments of need, desert, likelihood of success etc. themselves. To illustrate: A sexist is very unlikely to first correctly register that a female and a male party are equally deserving, and only then make a partial choice. He will fail to see the female party as equally deserving in the first place. And where he grudgingly admits that parties are "equally deserving" etc., we have reason to suspect that he downplays certain differences that actually weigh in favor of the female party.¹⁴

¹⁴ Could we reply that judgments of need and desert at least have some basis in observable facts, so that we have some reliable standards that can help us to detect possible biases? If the defender of lotteries is tempted to make

Thus, the decision to hold a lottery seems to be no less vulnerable to problematic biases than the decision to simply pick. Therefore it is hard to see how the general suspicion we are considering can make the decision to hold a lottery morally preferable to the decision to pick. Far from being morally required, lotteries start to look as problematic as picking is supposed to be.

Proponents of PA might reply that this is reason to use lotteries more extensively. If we cannot trust our judgments of need, desert, etc., we should use lotteries to bypass our reasoning capacities in more cases. Though this may sometimes lead to worse decisions, we cannot trust our ability to tell which decisions are better or worse. Better to use means that make decisions "based on no reasons" even where this may keep out good reasons that would make a difference.

Surely this proposal goes too far. It would lead to the requirement to hold a lottery in *all* allocation conflicts, no matter what the correct judgment of desert, need, etc. seems to us to be. This consequence shows how radical the suspicion invoked above is. Any suspicion that is general enough to lead to the Lottery Requirement will lead us much further – to a principle that puts our decision-making capacity as a whole into question. If defenders of the Lottery Requirement want to avoid such extremes, they have to put trust in our abilities to register the symmetry of conflict cases. But once they do that, it is unclear how they can maintain a general distrust in our ability to pick in these cases. For it seems that to thoroughly appreciate the symmetry of conflict cases *just is* to be in the right frame of mind for picking. As I said, I do think that the required trust is often warranted. What I have argued now is that defenders of the Lottery Requirement should agree, on pain of radical consequences.

To sum up: I argued that there are many other feasible procedures that are permitted by the norm on which most proponents of PA ground their defense. So lotteries are not necessary to satisfy this norm and hence are not required by it. This, as I said, is my main objection to PA.

But there is a reply we need to discuss. It is this: Maybe all the procedures I have suggested as putative alternatives to lotteries actually *are* lotteries. The one thing these

. .

this reply, then this time it is surely he or she who is being naïve. For instance, centuries of racism have surely proven that people keep judging objective features like intelligence etc. in massively distorted ways without in any way being bothered by evidence to the contrary.

¹⁵ Defenders of the Lottery Requirement might try to develop more complex solutions. They may devise other mechanisms that take on the task of registering the moral symmetry of the cases, and reserve the use of lotteries for a second step. We would have to know what these mechanisms look like in order to discuss this proposal. Especially, we would have to know how we can be sure that they are adequate without giving priority to our own judgments after all. But, more importantly, we should reconsider whether our grounds for a general suspicion are good enough to lead us to work towards making our reasoning capacities superfluous in all these cases.

procedures have in common is that they are based on arbitrary factors. If I go by who blinks next, or by dates of birth, or if I let my whim decide, do I not exploit a lottery mechanism (of an outer or inner kind)?

Let us call this the Inclusivist Reply. To assess it, we need to be a bit clearer about lotteries. It will not be necessary for our discussion to give a definition, but I will at least say this: In a narrow sense, a defining feature of a fair lottery is equiprobability. In this sense, a fair lottery must allow us to say that due to it, there was an equal chance for both A and B to be saved. As far as the interpretation of this equiprobability criterion is concerned, I will cut the proponent of the Inclusivist Reply some slack and be very liberal. One rather demanding way for a lottery to give equal chances would be to yield equal objective chances. 16 The next best thing would be a process that is objectively and symmetrically counterfactually open (in the sense defined in Hare 2012, 380 ff). The toss of a fair coin might be such a mechanism. However, I will also include mechanisms that generate equiprobability of an epistemic kind on the part of the rescuer. If you have a choice set-up that is designed to be perfectly symmetrical, so that, whatever you do, your reasons to believe that A will be saved and your reasons to believe that B will be saved are exactly similar, then that process generates equal epistemic probabilities. Pulling straws may be such a mechanism. I will count that as meeting the equiprobability criterion – although, as I said above, some of our intuitions call for more objective chances (cf. the case in which you falsely believe your coin is not biased).

Let us call all mechanisms that generate equiprobability in some such sense fair lotteries in the narrow sense. Then there are two ways to understand the above proposal (i.e. the Inclusivist Reply). On the one hand, it may be the view that the procedures I discussed yield equal chances in a relevant sense, i.e. that they are fair lotteries in the narrow sense. On the other hand, it may purport to use a sense of "lottery" in which fair lotteries need not give equal chances of *any* kind.

How could a defender of PA argue that even the procedures considered last give the candidates equal chances? On the face of it, this task looks hard, because the feasible mechanisms I have discussed do not even require the belief that there are equal chances. If you simply pick, or if you go by dates of birth (maybe in the way suggested by Kornhauser

¹⁶ To offer and defend an account of the appropriate notion of *objective chance* would be the task of the defender of the Lottery Requirement. But there are several ways one might go. If one is an indeterminist, one might claim that tossing a coin is a genuinely indeterministic process. Lewis (1986, 84 and 117 ff) suggests this, although he admits that the objective chances involved will hardly be exactly 50:50. If, by contrast, one is a determinist, there might still be the option of accepting a notion of objective *deterministic chances*, Schaffer's (2007) objections notwithstanding. Thus one may claim that the outcomes of the toss of a fair coin have an objective 50% chance in deterministic worlds. Indeed, one may follow a suggestion by Eagle (2011) and claim that tosses of fair coins are among the paradigms that implicitly define the chance role.

and Sager), how can there be *any* sense in which you can truly say that both parties got equal chances in the choice process?

To begin, it would surely be hopeless to point to a past history of chance factors. Suppose we were told this: "Look, the history of events that led to your picking is full of factors that are just as chancy as coin tosses are. Even if, *at present*, your picking does not make both outcomes equally likely, there were times in the past at which chance factors made them equally likely. If we look at the whole history, it is a lottery after all!" This is a non-starter. First, it is not clear why we should be entitled to assume that there ever was a time at which both outcomes were rendered strictly equally likely by chance factors. And *if* we are entitled to such a speculative assumption, it is not clear why we are not entitled to assume the same with regard to *every* choice that was ever made. Surely, we have cast our net too wide. The Lottery Requirement would be trivialized.

More importantly, the mere dependence on a chance factor in the past does not intuitively suffice to do the job of a lottery. Take an example of a conflict case involving a lifeboat. The captain is a racist who lets someone die because of the color of their skin. But a couple of days before, the captain's colleague tossed a coin to decide whether to accompany the captain on the present shift. If the coin had landed differently, the colleague would have been present and the racist captain would have chosen differently (in order to conceal his racist views from his colleague, who is already suspicious). So there was an earlier point at which a chance event made a different decision equally likely. But it seems that the presence of such a factor does not make it the case that the decision is impartial and gives equal chances in the intuitively appropriate way.¹⁷

Defenders of PA might instead choose to appeal to the decision-theoretic principle of indifference (also known as the principle of insufficient reason). It says that if you have no idea which outcome is more likely, you should consider all of them *equally* likely. The defender of PA might say that all the procedures mentioned before give you no reason to consider it more likely that A will be saved than that B will be saved and *vice versa*. So all these procedures should be considered equiprobable. But, first, there are well-known general problems with this principle.¹⁸ Secondly, some of the cases above raise a special problem. As

_

¹⁷ I am not quite sure what is missing in this case. But if I had to speculate, my guess would be this: For a chance factor to constitute a lottery, it must be used *with the intention* to make the decision depend on this chance factor. ¹⁸ The standard objections are familiar from textbooks on decision theory. A technical worry is that states of nature can be parsed and re-parsed in different ways, yielding incompatible assignments of equal probabilities (see, e.g., Luce and Raiffa 1957, 284f). A more philosophical worry is this: If I have no idea how likely the outcomes are, I have no reason to treat them as equally likely any more than I have reason to treat them as unequally likely (see Resnik 1987, 37).

Lewis (1986) has argued, beliefs about objective chances constrain rational credence (cf. his *Principal Principle*). In some of our cases, you have *very* strong reason to believe that the chances of the outcomes, whatever they are, are *not* equal (cf. Kornhauser and Sager's method). This is in tension with the claim that you would be rational in assigning .5 subjective probabilities to both of the outcomes. (Instead, I think you should refuse to assign any definite subjective probabilities at all – pretty much as in decisions under ignorance.) The same holds for many acts of simple picking. If I simply pick one of the parties, it is quite likely that able psychologists could tell a story about why, given the patterns in my attention etc., it was very probable that this person would be picked. I argued above that this does not render these processes objectionable. But it does mean that I will rarely, if ever, be justified in pretending that all pickings I might have made were strictly equally likely.

In view of these worries, we will turn to the second interpretation of the Inclusivist Reply. Proponents of PA might claim that even though the processes we have considered do not give the candidates equal chances in any sense, they can still count as fair lotteries in a wider sense. What is crucial about all of the procedures above, they might say, is that they involve some element of arbitrariness, and that this is really all they mean when they appeal to a duty to hold "a lottery".

One problem with this move is that the step from genuine equiprobability to arbitrariness is bigger than the proponents of PA will want to admit. I am not going to argue over whether the term "lottery" is apt to cover all of the procedures mentioned above (which even involve arbitrary picking). The crucial point is that the step from equiprobability to arbitrariness seems to lose much of what is intuitively compelling about the Lottery Requirement. If we employ an equiprobable chance mechanism in a conflict case, we can later say to the losing party that we took care to ensure that he or she had the same chance to be saved. If our choice was merely arbitrary, all we can say to the loser is that we decided to save the other party for no particular reason. This sounds considerably less consoling. I take it that this is part of the reason why the proponents of the other accounts (SSA and ICA, to be discussed below) will not consider all of these arbitrary procedures lotteries (or permissible, for that matter). And insofar as I share the intuitions behind the Lottery Requirement, my intuitions require equal chances, not mere arbitrariness. As a procedural moral standard, arbitrariness is much less intuitively compelling.

This leads to a second problem. It can be dubbed, somewhat grandiosely, the collapse into epiphenomenalism. The root of the problem is this: Choosing arbitrarily is not a particular *way* to make decisions "based on no reasons" (as Stone put it). To say that a choice

is arbitrary is just to *describe that* it is (somehow) based on no reasons. More specifically, we cannot see the arbitrariness shared by the above procedures as a means to keep out bad reasons. In conflict cases, arbitrariness is simply what results when they are kept out. Consider: We stipulated that in conflict cases, there are no *good* reasons to save one party rather than another. So if we rule out the *bad* reasons as well, of course all that remains will be a choice that does not go by *any* reasons and that will thus be arbitrary. But we have no reason to believe that this arbitrariness is what *makes* these procedures permissible. They somehow keep out bad reasons, but being arbitrary is not what keeps them out. Now, given that (as I have pointed out above) it is not terribly plausible to see arbitrariness as a moral advantage anyway, what reason do we have to think that this arbitrariness is itself required, rather than a mere consequence of something else (i.e., avoiding bad reasons) that is required?

As long as we thought of lotteries as special mechanisms to generate equal probabilities in some sense, we could argue that these lotteries are required as ways to keep out bad reasons. But my discussion has suggested that there are permissible procedures (such as picking) that do not involve any such mechanism. Their arbitrariness is not a randomizing mechanism, so we are left without a clear reason for believing that it is what makes these procedures required. The fact that all permissible procedures are arbitrary is fully explained in terms of a requirement that has nothing essentially to do with arbitrariness. What is required is that a choice not be based on bad reasons – and in conflict cases (due to the absence of good reasons), it just so happens that all of these choices are arbitrary. Arbitrariness seems to do no moral work; it is an epiphenomenon.

But what of it? Even if the Lottery Requirement merely picks out a moral epiphenomenon, could it not still be true? I do not think so. Statements about what is morally required seem to generate intensional contexts. If we say that it is required in all cases of type X that A be F, then being F had better be something that is itself required, not just any property that is coextensive with something that is required in all cases of type X. Compare: Due to the structure of conflict cases, any permissible course of action will save no more than one party. But surely we should not defend the principle that in conflict cases, it is required to save no more than one party. It just so happens that these are the only permissible courses of

-

¹⁹ It may be obvious, but let me stress that *not being based on bad reasons* and *being arbitrary* are different properties for decisions to have. They are not even coextensive; many decisions avoid bad reasons without being arbitrary (namely, those that are based only on *good* reasons). These properties merely *coincide* in conflict cases (due to the lack of good reasons). Being arbitrary *entails* not being based on bad reasons, but it is not a *means* or *way* to avoid bad reasons.

action in these cases; it would be false to suggest that saving no more than one person is itself something that is morally required.

This is what I have called the collapse into epiphenomenalism. To wrap up this discussion, here is the bottom line: I argued that there are a huge number of choice procedures that are procedurally fair, as required by PA. But they cannot all be counted as lotteries in a narrow sense that treats some kind of equiprobability as criterial. However, if we adopt a broader notion of lotteries as merely arbitrary choice processes, we no longer have in hands a feature that is morally required as a way to keep out bad reasons. Thus, both ways to interpret the Inclusivist Reply fail.

However, we are still not finished with PA. Maybe we went wrong in our initial choice of a procedural norm. Maybe, that is, the relevant procedural norm that leads us to lotteries is not, or not *just* a norm of impartiality. For example, one may appeal to a norm of *transparency*. Even if there are other ways to keep out bad reasons where necessary, a lottery can assure all people concerned that no bad reasons are in play (cf. Wasserman 1996). It is a choice procedure that not only avoids racism, sexism, etc. but also makes it public and transparent that they play no role.

With regard to decisions by private individuals, I do not think such a norm of transparency is valid. Granted that we must not act on bad reasons. But are we generally obliged, in addition, to take positive measures to prove to others that our hearts are pure? Will any action of ours be wrong if it fails to positively assure others of our good motives? This is doubtful, to say the least. (Note that this is not just the familiar requirement that we should be able to justify our actions if people should ask for it. It is the requirement to act only in ways that are sufficient by themselves to prove that no bad reasons could be present.)

Still, there might be other cases in which such a moral requirement of transparency is valid. One example may be public institutions. Many such institutions live on public funding, and some of them have coercive power. This may put them under a requirement to make their guiding principles visible to everyone. Many examples of conflict cases will occur in such institutional contexts (e.g., organ allocations).²⁰ In the context of such institutions, a randomizing mechanism could *serve as a means* to achieve transparency. (So the epiphenomenalism worry would not apply.)

Still, I do not think that randomizing mechanisms are the proper way to make sure and transparent that institutions are impartial. In a nutshell, the reason is this: Institutions are

-

²⁰ Thanks to an anonymous reviewer for pressing this point.

subject to special criteria of fairness and impartiality, criteria that do not apply in the case of individual one-off decisions. We judge institutions (among other things) in terms of the long-term, society-wide distributive effects of their rules and procedures. As I will argue, these distributive criteria make it hard to see what role a lottery can play. My argument to this effect proceeds in several steps. I start with the following example:

We design an institution that deals with a certain type of conflict case. We want to rule out sexist biases by means of a lottery.²¹ One option is to run a new lottery in each conflict case in which we have to decide between a male and a female party. However, the outcomes of such a single-case lottery are likely to mirror the proportion of men and women in the population only if the number of conflict cases is very large. So we would run a risk that our institution's official coin turns out to select, say, disproportionately many men. But we have an alternative. We can pursue a perfectly equitable distribution and give all parties equal chances with the following mechanism. We first construct an algorithm for future choices that respects the distributive aim. If the aim is an equal distribution, it can be a simple pattern: FMFMF... ("On the first occasion, save the female. Then the male, then the female, then...") Next, we construct a permutation of this series – in our case, the reverse pattern MFMFM... will do. Then we let a fair coin decide between the two series. Firstly, this way all future candidates will have received equal chances. (Whenever the chosen series will dictate a particular choice in the future, we can say that we have used a mechanism that made the opposite choice equally likely.) And, secondly, the procedure guarantees a fair distributive pattern.

Call this mechanism the Lottery Between Algorithms. I think that upon reflection, it can be seen to be superior to a single-case lottery. As far as transparency and impartiality are concerned (which is what proponents of PA appeal to), it is hard to see how a single-case lottery could have advantages. And to simply assume that we owe each party a new lottery would beg the question.

One objection might be that the Lottery Between Algorithms forces officials to constantly keep considerations of gender in mind, and to base decisions on them. A single-case lottery, by contrast, would make it superfluous to think in these categories at all, making it a more effective tool to eventually overcome sexism. But this objection misses something. The problem of institutional sexism (or racism, etc.) is not just that people think in terms of sex (or race, etc.). It is also that our institutions can have manifest distributive effects that are

_

²¹ We may stipulate that in this imaginary case, our society is very homogeneous apart from differences in sex, so there is no reasonable fear of any other kind of bias.

inequitable. I think that correcting the latter flaw is more urgent (and more realistic) as a goal for social policy. Consider: Suppose our coin has so far saved mostly men, and we all think it is time that a woman should be saved for a change. But, alas, the coin again selects a man on the next occasion. How can we justify this result, given that we had available a mechanism that would have achieved equal chances *and* an equitable distribution? (Again, it would beg the question to just assume that we owe individuals a single-case lottery.) Can we seriously maintain that it was more important to ensure that the individual officials do not pay attention to gender? I think this would be an inappropriate concern with individuals' motivational hygiene. It is more important to design institutional rules that are explicitly committed to keeping distributive inequalities from arising. And since the Lottery Between Algorithms gives us both that and equal chances, I think we should prefer it.

Some might still feel uneasy about the particular use of gender categories here. Even if distributional aims do have great importance, one might think that the Lottery Between Algorithms forces us to use gender categories in a way that is objectionable. It would lead to a problematic kind of affirmative action in which we will often fail to save a certain party *because* he is male (say). This, one might say, will be discriminating. A single-case lottery would avoid this.

I think that this type of objection to affirmative action is confused.²² As Anderson (2010, 155ff) convincingly argues with respect to race, the relevant group categories play different roles in discrimination and in affirmative action. As applied to our case: A sexist will think that the male candidate, because he is male, has more of the morally relevant characteristics, e.g. that he is more deserving. Nothing like that is implied by the Lottery Between Algorithms. By using it, we do not treat gender as bearing on the *worthiness* of being saved. Here it is crucial to remember that the idea is to use such a mechanism in conflict cases (in the sense defined above), i.e. in morally symmetrical cases. It is only because these are cases in which it is clear that the worthiness of individuals does not settle the issue either way that we allow *additional* considerations (such as large-scale distributive aims) to play a role. These further concerns are not treated as considerations that outweigh or replace questions of individual desert, or need, etc.

So, again, I think that proponents of PA must agree that the Lottery Between Algorithms is to be preferred over a single-case lottery. But here, finally, is the rub: If we

²² Let me also note that another typical objection to affirmative action does not apply in our cases. Affirmative action is often criticized for not achieving certain aims (e.g., a higher number of successful university graduates from relevant groups, cf. Sander 2004). In our cases, there are no such independent criteria of success. Here our only aim is to save persons without giving disproportionate advantages to some groups.

choose it, I find it hard to see *why the lottery was needed in the first place* – unless we beg the question by assuming that giving everyone equal chances matters as such. If all we want is to make transparent that our procedure is impartial (and this is all that proponents of PA are entitled to appeal to here), it does not seem to matter whether some person simply picked an algorithm, or whether a coin did. If a person picked it, that cannot possibly count as a way to be partial towards either group. (If you were a sexist, which of the two series would you choose? How could it matter to you?)

So it seems that as far as impartiality and transparency are concerned, nothing significant is lost if we do not hold a lottery but let someone pick a series. The design of the series seems to transparently guarantee all the impartiality that one could ask for. But if *that* is true, we might as well take another step and introduce a *quota*: We let officials pick on each occasion as long as the overall result satisfies certain proportions. The sheer number of pickings that occur can hardly make a difference. And if everyone knows that for each male party saved, a female party has to be saved (and *vice versa*), this again rules out any suspicion of sexist partiality. (If you were a sexist, would you save more males first and then more females later, or the other way round?)

Thus, I do not think that this step-wise progression from the Lottery Between Algorithms to a quota reintroduces objectionable forms of discrimination. Officials are bound by a quota as much as they are by the chosen algorithm. They are thereby forced to choose in a way that is conscious of gender categories, but neither the quota nor the algorithm treat gender as a consideration that is supposed to outweigh or to replace considerations of individuals' desert etc. (Of course, sexist officials may still enjoy saving people of a particular sex more. But this will be just as true if we use a single-case lottery instead. All that should matter is whether a procedure leaves officials any room to further the gratification of these emotions. And a quota seems to rule this out as well. Again, a desire to save men rather than women is no more gratified by starting to save more men first and more women later than it is by the reverse series, or by tossing coins.)

The following result emerges: Institutional impartiality is a matter of distributive aims, and the more we do to realize them, the more it seems that the use of chance mechanisms is at best superfluous. I have considered a simple case in which a perfect combination of equal chances and distributive fairness was possible, and I concluded that such a procedure has no advantages over a more direct approach which involves a binding quota that sets the limits within which the officials may freely pick. Of course, in reality a Lottery Between Algorithms will often not be feasible. But the use of quotas does seem to be possible. (Indeed, in very

many cases it is *actual*.) And if they are at least as good as the *perfect* combination of equal chances and distributive fairness that was available in my easy case, they can hardly be inferior in more complex real-life cases in which combinations of equal chances and distributive fairness will require compromises.

In sum, norms of transparency cannot ground the Lottery Requirement with regard to institutions. For institutional contexts, my objection is not just that there are alternative means that are as good as lotteries. Rather, those alternative means seem to me to do an even better job.

A final move we may consider departs from the "keeping out bad reasons"-line altogether. Maybe what matters is not whether you decide on bad reasons or not, or whether you just pick. Maybe what matters is that people should not decide between people *at all*. Maybe there is hubris in any such choice, no matter on what grounds it is made. I think there is something to this view, but it does not give us any idea why there should be less hubris in declaring that a coin is to decide between lives. It most certainly feels wrong to rule that a person will have to die. But shouldn't it feel just as wrong to rule that a person has to die if one's coin shows heads? Our choice will make all the difference between life and death anyway, no matter whether the choice is direct or whether it is an assignment of acts to outcomes. We should not close our eyes and pretend that the coin does all the choosing here. So although this move correctly registers the great weight of the choice we face, it does not show how tossing a coin is more than false relief.

3. The Ideal Consent Account

The findings so far suggest that lotteries are not required by fairness, either in a distributive or a procedural sense. But there is still another approach that defenders of the Lottery Requirement might take. They might appeal to a form of ideal consent or agreement, by arguing that holding a lottery is the unique procedure that everyone concerned would have sufficient reason to accept.

There are actually two quite different reasons why such an approach may seem attractive. First, some contractualists may endorse this approach just because they think that the question of what all parties concerned have sufficient reason to accept is generally at the heart of morality (a general view famously proposed in Scanlon 1998). But, second, there is a different reason to take this approach seriously, a reason that needs no heavy-duty theoretical superstructure. Consent or agreement may matter because they can turn the lottery procedure into a consensual gamble, and such a gamble has characteristic normative powers that may

ground a novel defense of the Lottery Requirement. To put the point briefly, the previous accounts have tried to show that a lottery is required to satisfy relevant claims of the parties in need, either claims to the contested good or claims to a fair decision process. However, what if a lottery is required, not to *satisfy* the parties' claims but to *change* them?

Consider this example: You and I argue about who gets to use the car this afternoon. We both have equal claims. Eventually, we agree to let a coin decide. If you win and I lose, I no longer have an equal claim to take the car. You are entitled to take it, and if I protest, I violate an agreement. In a case like this, a lottery procedure has the status of a consensual gamble. As such, it has normative power: it can change our entitlements. This, as I said, opens up a new possibility. Maybe the defender of the Lottery Requirement can reason as follows: At present, A's and B's conflicting claims seem to leave you no permissible choice. (After all, we are assuming that without a lottery, it is not permissible to save either party. Nor, of course, is it permissible not to save anyone.) But a lottery, considered as a consensual gamble, allows you to change these claims in such a way that is becomes permissible not to save the loser. Thus, a gamble would be the only permissible way out, and so morally required. This is a third defense of the Lottery Requirement.

In what follows, I will mainly discuss the second version of ICA (i.e., the one that appeals to the normative effects of a consensual gamble, not the general contractualist one). But my points will apply to both versions. An important ingredient of ICA must indeed be the consent of the parties. It is absurd to suppose that the actual entitlements of persons could be changed by the toss of a coin without their consent. (If we both have an equal claim to take the car today, you cannot just pull out a coin, shout "heads", toss it and, if you are lucky, rush off.) Of course, defenders of the Lottery Requirement should not appeal to *actual* consent. You will rarely be in a position to ask, say, two castaways on the open sea whether they consent to a lottery. Nor should we speculate about whether the parties would, as a matter of fact, consent if we could ask them. (I suspect that each would insist that he or she be saved, period.) But maybe we can argue that in all conflict cases, the parties should offer consent in some normative sense?

D. Estlund (2008) argues that political authority can be based on, not actual consent but *normative consent*, where normative consent is understood as "the moral duty to offer consent when given the chance" (ibid., 134).²³ The presence of such a moral duty, Estlund claims, can constitute a nullifying condition for non-consent. Very roughly, this means that the fact that

-

²³ Thanks to an anonymous reviewer for suggesting that I refer to Estlund's views in this context.

someone morally should consent or should have consented can play part of the justifying role of his or her actual consent. Maybe we can apply Estlund's idea to our case. Maybe, if both persons morally ought to consent, that is sufficient to endow a lottery with the required normative power.

I do not think there is much hope for this proposal. If there are moral reasons *to consent* to a lottery procedure, it is very natural to expect them to derive from the moral reasons there are *to hold* a lottery. If it is not otherwise obligatory to hold a lottery in conflict cases, why would it be morally required for the parties to *consent* to such a procedure? And in fact, when I try to think of reasons to consent to a lottery, I think of things like fairness and the equal claims of each of the parties, all of which would primarily weigh in favor of a lottery, and only derivatively in favor of consenting to a lottery. But if this gets the priorities right, the proposal will not help. Throughout this paper I have been concerned to find such reasons to hold lotteries. And I have argued that considerations of fairness etc. fail in this regard. So we cannot point to them as reasons to consent to lotteries either. More generally, the attempt to argue that we are morally obliged to consent to a lottery risks begging the question by presupposing what we are trying to decide here.

Are there moral reasons for the parties to consent to a lottery that do not derive from moral reasons to hold a lottery? I cannot disprove this idea, but I find no promising candidates.

But maybe we can introduce a *non-moral* version of normative consent. Maybe, that is, it is enough that the parties simply ought (in a non-moral sense) to consent. I am not sure about this, but let us give this proposal a hearing. Often, the reasons we have to let a coin decide are pragmatic. Compare again the above example: You and I argue about who gets the car, and tossing a coin is the single most reasonable way to resolve our conflict. If, in such a case, I irrationally refuse to toss coins, maybe my refusal is null. Maybe you are within your rights to just shout "heads", flip a coin and treat the result as binding no matter what I say. Again, I am not sure that such 'normative consent' would indeed suffice. But for the sake of argument, let us grant it. Do the parties in conflict cases have such reasons? This is unclear. If you and I argue about the car, a lottery may be the best way to resolve the conflict. But that is because none of us backs down, and the option of using physical force is less reasonable. Such a pragmatic reason will often be absent in conflict cases. If you alone must decide where to steer your lifeboat, nobody keeps you from just picking. So what other reasons would the parties have to consent?

To find other non-moral reasons to consent, we may ask: Under what condition should the parties think that they would *benefit* from a gamble? A natural answer would be this: They should do so if they fear that they might otherwise have an even lesser chance. In that case, they might well consent to the lottery in order to commit the rescuer to at least give them an equal fifty-fifty chance. Taking this line would have a major advantage. Unlike defenders of PA, we can explain why it is precisely an equiprobable gamble that the parties should consent to. By assumption, both parties believe that they stand to gain from an agreement. This turns the situation from a pure conflict into a situation where bargaining is rational. And since the situation is symmetric, all of the major views on rational bargaining (the Nash approach, the Kalai-Smorodinsky approach) will lead the parties to converge on mixed strategy of 50-50, that is, on a gamble with equal odds. (Of course, one might ask why we should regard the situation as symmetric. Maybe one of the parties is less afraid of death and would use this to increase his or her bargaining power. But, for the sake of argument, I give the present approach some credit and grant this assumption as well.)

Still, this proposal fails. First of all, it is not clear why the parties in such cases should generally have the suspicion that the initial odds are against them. But this assumption is crucial; bargaining is rationally required only if it will improve your position compared to your fall-back position. More importantly, the reason we are considering will ruin the legitimacy of the agreement. The proposal is that the persons entitle the rescuer to give them a fifty-fifty chance because they fear they have a lesser chance otherwise. But, as things stand, this is the fear that there is a high chance of *being wronged*. (Again, defenders of the Lottery Requirement hold that it is wrong to save a party before a lottery is held.) If the parties waive their claims under that condition, that hardly legitimizes the rescuer's action. It would be like "allowing" someone to beat you up because you fear that he might do worse things otherwise. The parties' consent would be given under duress.

Another possibility would be that the parties in need have an ambiguity aversion (in the sense revealed by the Ellsberg paradox, see Ellsberg 1961).²⁴ So they do not necessarily fear that they have a lesser chance without a lottery. They merely have a preference for knowing their odds instead of having to deal with uncertainty, especially where so much is at stake. While I admit that I myself might be prone to ambiguity aversion, I am not sure that all rational agents are. Indeed, the Ellsberg paradox is usually called a paradox because it is hard to square with the minimal rationality assumptions underlying expected utility theory.

-

²⁴ Thanks to Wolfgang Spohn for bringing up this possibility.

Anyway, *if* ambiguitiy aversion is rational at all, it is at best rationally *permissible*. This is not enough to yield normative consent.

I see no reason to suppose that people in conflict cases generally have reasons to waive their unconditional claims and to enter an agreement to the effect that a lottery should decide who will survive. ICA fails just like SSA and PA. We found no moral reason to hold a lottery.

4. Arguments from Persisting Intuitions and from Moral Uncertainty, and the Prospects for a Weaker Principle

I have discussed three general ways to defend the Lottery Requirement. I have concluded that none succeeds. This suggests that skepticism towards the Lottery Requirement is warranted.

But even if my arguments are convincing, there are at least two remaining grounds for resistance to this skeptical conclusion. I call them The Argument from Persisting Intuitions and The Argument from Moral Uncertainty. The first starts from the fact that the intuitive appeal of the Lottery Requirement persists. I suspect that it will still *seem* to many readers that a lottery is obligatory in conflict cases. I confess to sharing this feeling. But let me be clear right away that it would be implausible to claim that we know the truth of the Lottery Requirement directly on the basis of this intuition. Whatever we think about moral intuitionism, the Lottery Requirement is an unlikely candidate for a fundamental principle the truth of which could be directly intuited.

But the Argument from Persisting Intuitions may take a different route. While our shared intuitions may not directly entitle us to believe the Lottery Requirement, they may indirectly cast doubt on my case against it. They may be evidence that I missed something, that there must be good reasons behind the Lottery Requirement, even though we may not have identified them yet. "After all," one might ask, "how could so many of us, and so many philosophers, be so wrong?"

This argument is unconvincing. Firstly, while I share the relevant intuitions to some extent, I also think that they do not go very far. Ask yourself: If, in a conflict case, the captain of the lifeboat happens not to have a coin with him, and if there are also no straws or wheels of fortune available, is he doomed to be immoral? Would we complain if he just picks? How could we? Could this mean that there *is* a requirement to hold a lottery, but that we can easily escape the scope of this requirement by failing to bring an adequate lottery device? This would seem to be a very undemanding requirement indeed. Or maybe what is immoral in these cases is the failure to bring a suitable coin or wheel? I have never heard anyone utter a moral complaint about that.

So while I agree that our intuitions persist, I am not sure how firm they really are. Secondly and more importantly, these intuitions are not trustworthy, for debunking explanations come to mind too easily. First, as I argued in § 1, it is tempting to take too literally talk of an "equal distribution of chances", as if we were really handing out shares of some good. Secondly, as we saw in § 2, it is tempting to infer, from the fact that all adequate procedures in conflict cases have an arbitrary character, that this arbitrariness is what makes the procedures morally adequate. But this would be a fallacy. Thirdly, "chance" is ambiguous. When I say: "Everybody should have the chance to live a decent life", I do not use a probabilistic notion. I mean that everyone should be able to live a decent life, not (just) that everyone should be *likely* to do so. Thus, many requirements of "equal chances" could be less misleadingly described as requirements of equal opportunity. In consequence, it is tempting to support any so-called requirement of "equal chances" just because one is generally in favor of requirements of equal opportunity. Finally, our intuition that in conflict cases, a lottery makes a moral difference may be grounded in the wish to believe that there is something we can do to make very tragic cases less tragic. These temptations explain the persistence of our intuitions – but they also clearly undermine their status as evidence.

A second argument is the Argument from Moral Uncertainty.²⁵ Nothing I have said entails that it is wrong to hold a lottery in conflict cases. But, as I said, many of us have an intuition to the effect that it is wrong not to hold one. So one might argue that morality requires us to err on the side of caution. According this view, if we are definitely allowed to do something and there is a chance that it might turn out to be wrong not to do it, we are morally required to do it. Now, I think certain principles of moral uncertainty might hold in cases in which seemingly compelling arguments conflict (see, e.g., Lockhart 2000). But this is no such case. I examined the major arguments for the Lottery Requirement and found them all wanting. So at present, our only basis for resisting the skeptical conclusion is our persisting intuitions. And no principle should count any old intuition as meeting the standard for moral uncertainty, lest it render all our intuitions, no matter how unfounded (see above), credible by default. Absent better arguments to back these intuitions, there is no ground for moral caution. Also, while it is true that holding lotteries seems morally unproblematic, it does not follow that treating the Lottery Requirement as a valid principle is unproblematic. If we do so, we may end up blaming people for things (like merely picking people in conflict cases) that are not wrong at all. If moral caution is required, this point cuts both ways.

_

²⁵ Thanks to Daniel Ramöller for discussion here.

A final move I want to discuss is to admit that the Lottery Requirement as I have formulated it lacks support, but to insist that this was not the best formulation anyway. Maybe the thought is not that there is an obligation to hold a lottery. Maybe the idea is just that we *ought* to hold a lottery in a weaker sense that does not imply obligation – maybe it is just advisable, or the best course of action, without being required? Such a weaker principle (we should not call it a lottery *requirement*) may seem to find support in some of the arguments I rejected above. For example, we might say that the mere possibility of a hidden bias is a reason to hold a lottery, even if, as I argued, it does not suffice to ground an obligation to hold one. The same might go for our persistent intuitions.

I will not try to show that no such weaker principle can be defended. But it is hard to see what can be said in its support. Take the example of the possibility of a hidden problematic bias. As far as I can see, we can raise an exact analogue of the problem discussed above for weaker principles. Suppose we accept that the mere possibility of such a bias is enough, in principle, to make it the case that we ought (in some sense) not to pick in conflict cases (although we may still do it). If that is the case, it seems that this possibility will *also* make it the case that we ought not to rely on our judgments about the conditions for a lottery (e.g., equality of desert, need, etc.). Again, since the decision to hold a lottery seems to be just as vulnerable to problematic hidden biases, it is hard to see how the possibility of such a hidden bias can make this decision advisable.

But maybe other considerations we have discussed can still be used to support a weaker principle? I am skeptical. In general, it is a mistake to think that a weaker principle can automatically derive support from all the same intuitions and arguments that (putatively) support an obligation to hold a lottery. The considerations discussed above were mostly intended to show that holding lotteries is a requirement of fairness (in some sense or other). Fairness is a paradigm consideration that renders actions not just advisable but *obligatory*; even frameworks that treat it as a *pro tanto* principle treat it as a *pro tanto duty*. So to the extent that the considerations in question are convincing, they will support just such a duty. To the extent that they are not convincing, they are – well, not convincing. At any rate, it is not clear that they can still serve to show that lotteries have a quite different normative status that has nothing to do with fairness or the like.

Much the same, I think, goes for our persistent intuitions about lotteries. I have already admitted that I share these intuitions to some extent (though I have also highlighted their limits). But the intuitions that I do have are not adequately captured by a principle that treats lotteries as advisable but not obligatory. For instance, I keep having the feeling that if we fail

to hold lotteries, this is unfair and the parties in need have a legitimate complaint. According to the principle we consider now, this intuition is misguided, just as it would be if lotteries were morally neutral. If we must let go of these central intuitions anyway, are we sure there is significant support left?

It is not an accident that weaker lottery principles are not the principles that are actually discussed in the relevant literature. As far as I am aware, all of the authors who support lotteries formulate their support in terms of fairness, justice, entitlements, claims etc. (see the references in footnote 3). And, again, these are just not the types of considerations that could give us advisability or goodness without duties (either *pro tanto* or all things considered). So, to sum up, it is not clear what can be said in support of a weaker lottery principle or who actually endorses it.

5. Conclusion and Consequences

I have already hinted at the implications of my skeptical conclusion for our practice. On the one hand, nothing I said suggests that we must not hold lotteries to allocate goods in conflict cases. In fact, there are often good non-moral reasons to do so. There are pragmatic reasons: It can be an effective and fast allocation procedure and avoid long arguments. There are psychological reasons: It can relieve the burden of having to make certain hard choices. On the other hand, however, absent better defenses of the Lottery Requirement we are *not* entitled to treat this alleged requirement as valid, e.g. by using it in our practices of moral assessment and sanctioning.

In addition, my skeptical view has implications for a range of philosophical debates and arguments. In closing, let me mention some of these implications. To begin, an alleged duty to hold a lottery has received discussion in the debate sparked by J. Taurek (1977). Famously, Taurek discusses what we may call *unequal-number cases*: You can save either a smaller group or a distinct larger group, but not both. Again, all else is equal. Taurek held that in such cases, you are not required to save the greater number. What, then, should you do? Taurek says what he *would* do is toss a coin to give everyone an equal chance (cf. 1977, 303). Importantly, he does not say that this is what one *should* do. Indeed, there are reasons for thinking that his view cannot have been that one is required to hold a lottery.²⁶ Anyway, most

²⁶ In fact, one main argument in the first part of Taurek's paper goes as follows: If you "know and like" the one person in need (David), you are permitted to save him rather than a group of five strangers. But this, Taurek stresses, is *not* the kind of special relation that can override an alleged duty to save the greater number: "In securing David's survival I am acting on a purely personal preference. It is the absence of any moral requirement

of those who buy Taurek's argument against a duty to save the many hold a different view. They think it is a requirement of fairness to hold a lottery. Indeed, even those who do *not* buy Taurek's argument find this stronger verdict so obvious as a consequence that they ascribe it to Taurek.²⁷ Some even suggest that Taurek's objection against a duty to save the many *is* precisely that we owe every person an equal chance.²⁸

My argument counts against these latter views (but, as I said, it may not count against Taurek's actual view on tossing coins). *If* Taurek is right that we need not save the greater number, then unequal-number cases, too, are cases where equal claims conflict and all else is equal. So my arguments apply. If we follow Taurek in rejecting a duty to save the greater number, we must accept that it is permissible to save the smaller number *even without a lottery*.²⁹ Many will think that this counts against Taurek's claims, probably including some of Taurek's supporters.

There is a further consequence I want to mention. In a series of papers (collected in Harsanyi 1976), J. Harsanyi has proposed an axiomatic approach to utilitarianism and proved a much-discussed theorem. It says, roughly, that if (i) all individuals in a certain population have preferences that satisfy the von Neumann-Morgenstern coherence axioms, and if (ii) an observer has social preferences that likewise satisfy these coherence axioms plus a Pareto condition, then these social preferences can be represented by a utility function that gives weighted sums of the utilities of the individuals. In a famous criticism of this proof, P. Diamond (1967) has objected that von Neumann and Morgenstern's independence condition is not acceptable as a condition on social preferences. This condition entails that the observer is indifferent between, e.g., a lottery that gives each of A and B a 50% chance to some good and simply giving the good to A. Diamond argues that the observer should not be indifferent

_

to save these others rather than David that makes my doing so morally permissible" (1977, 297). But if it is permissible to choose on a purely personal preference, it cannot be required to hold a lottery instead.

²⁷ Here is a sample of authors who champion this reading of Taurek: Regan (1983, 298) writes that "what we should do, Taurek thinks, is flip a coin to decide who to save"; Otsuka (2000, 288) claims that Taurek says that "we should follow a principle that directs us to flip a fair coin to determine whom to save"; Meyer (2006, 142) says: "we have a duty to flip a coin and give each of the four an equal chance of survival (Taurek)"; Hirose (2001, 341) claims that Taurek "thinks we should flip a coin to give each a fair chance of being saved." Not everyone accepts this reading, however. Kamm (1993, 76) reads Taurek thus: "If we want to show equal concern for all six people – though perhaps we needn't – we should toss a coin, thereby giving each [...] an equal (50 percent) chance." Doggett (2009, 302) gets the view exactly right: "[Taurek says] he *would* flip a coin to figure out who to save. Nothing hangs on this and it is controversial, but I think Taurek's view is *not* that there is a moral requirement to flip a coin."

²⁸ Timmermann (2004, 106) is one example: "Everyone should [...] be given an equal chance of survival. [...] [F]or Taurek, the decisive consideration is the fairness of the decision process."

²⁹ Gaus (unpublished) has independently developed a powerful critique of the Lottery Requirement as applied to Taurek cases.

here because it seems preferable to "give[] B a fair shake" (1967, 766). A very similar criticism can be found, e.g., in Sen (1970, 145).

In response, Harsanyi (cf. 1976, 68 ff) has insisted that indifference is the right attitude for the observer to take, and that Diamond's two alternatives really are on a par. In other words, while Diamond and Sen appeal to the Lottery Requirement, Harsanyi rejects it. Even many of those who are sympathetic to Harsanyi have described this move as "biting the bullet" (see, e.g., Risse (2002, 572)). In its stead, they have proposed repairs (like construing the fair or unfair processes as part of the consequences; see Risse (2002) and Broome (1991)). The argument in this article suggests that Harsanyi's stance may be justified after all. As I have argued, it is indeed hard to see why we should be morally required to give B his "fair shake". (In cases involving divisible goods, I will be the first to agree that utilitarianism has serious problems with fair distributions. But, as I said, it is a mistake to think that there can be a fair distribution in conflict cases.) That does not mean that Harsanyi's theorem is true, or that his axioms are without fault. But it does suggest that opponents of axiomatic utilitarianism (among whom I count myself) will have to look for the fault elsewhere.

References

Anderson, Elizabeth. 2010. The Imperative of Integration. Princeton: Princeton University Press.

Broome, John. 1984. "Selecting People Randomly." Ethics 95: 38-55.

Broome, John. 1990. "Fairness." Proceedings of the Aristotelian Society 91: 87-102.

Broome, John. 1991. Weighing Goods. Oxford: Blackwell.

Broome, John. 1998. "Kamm on Fairness." Philosophy and Phenomenological Research 58: 955-961.

Childress, James F. 1970. "Who Shall Live When Not All Can Live?" Soundings 53: 339-355.

Daniels, Norman. 2012. "Reasonable Disagreement about Identified vs. Statistical Victims." *Hastings Center Report* 41: 35-45.

Diamond, Peter. 1967. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility: A Comment." *Journal of Political Economy* 75: 765-766.

Doggett, Tyler. 2009. "What Is Wrong With Kamm and Scanlon's Arguments Against Taurek." *Journal of Ethics and Social Philosophy* 3.

Eagle, Antony. 2011. "Deterministic Chance." Noûs 45: 269-299.

Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75: 643-669

Elster, Jon. 1989. Solomonic Judgments: Studies in the Limitations of Rationality. Cambridge: Cambridge University Press.

Estlund, David. 2008. Democratic Authority. A Philosophical Framework. Princeton: Princeton University

Press.

Gaus, Simon. Unpublished. "Random Selection and Arbitrary Choice in Taurek Cases." Unpublished manuscript, Humboldt University Berlin.

Goodwin, Barbara. 2005. Justice by Lottery, 2nd ed. Exeter: Imprint Academic.

Greely, Hank. 1977. "The Equality of Allocation by Lot." *Harvard Civil Rights-Civil Liberties Law Review* 12: 113-141.

Hampton, Jean. 1994. "The Failure of Expected Utility Theory as a Theory of Reason." *Economics and Philosophy* 10: 195–242.

Hare, Caspar. 2012. "Obligations to Merely Statistical People." The Journal of Philosophy 109: 378-390.

Harsanyi, John. 1976. Essays of Ethics, Social Behavior, and Scientific Explanation. Dordrecht/Boston: Reidel.

Hirose, Iwao. 2001. "Saving the Greater Number Without Combining Claims." *Analysis* 61: 341-342.

Kamm, Frances Myrna. 1993. Morality, Mortality, Volume One. New York: Oxford University Press.

Kornhauser, Lewis A. and Sager, Lawrence G. 1988. "Just Lotteries." *Social Science Information* 27: 483-516. Reprinted in *Lotteries in Public Life*. *A Reader*, P. Stone (ed.), Exeter: Imprint Academic, 2011, 133-160.

Lewis, David. 1986. "A Subjectivist's Guide to Objective Chance." in *Philosophical Papers Vol. II*, Oxford: Oxford University Press, 83-132.

Lockhart, Ted. 2000. Moral Uncertainty and its Consequences. Oxford: Oxford University Press.

Luce. R. Duncan and Raiffa, Howard. 1957. *Games and Decisions. Introduction and Critical Survey*. New York: Wiley.

Meyer, Kirsten. 2006. "How to Be Consistent without Saving the Greater Number." *Philosophy and Public Affairs* 34: 136-146.

Otsuka, Michael. 2000. "Scanlon and the Claims of the Many Versus the One." Analysis 60: 288-293.

Regan, Tom. 1983. The Case for Animal Rights. Berkeley and Los Angeles: University of California Press.

Resnik, Michael D. 1987. *Choices, An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press

Risse, Mathias. 2002. "Harsanyi's 'Utilitarian Theorem' and Utilitarianism." Noûs 36: 550-577.

Sander, Richard H. 2004. "A Systematic Analysis of Affirmative Action in American Law Schools." *Stanford Law Review* 57: 367-484.

Saunders, Ben. 2008. "The Equality of Lotteries." Philosophy 83: 359-372.

Schaffer, Jonathan. 2007. "Deterministic Chance?" British Journal for the Philosophy of Science 58: 113-140.

Scanlon, Thomas. 1998. What We Owe to Each Other. Cambridge, MA: Harvard University Press.

Sen, Amartya. 1970. Collective Choice and Social Welfare. San Francisco: Holden-Day.

Sher, George. 1980. "What Makes a Lottery Fair?" Noûs 14: 203-216.

Stone, Peter. 2011. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford: Oxford University Press.

Taurek, John. 1977. "Should the Numbers Count?" Philosophy and Public Affairs 6: 293-316.

Thaler, Richard and Sunstein, Cass. 2008. *Nudge. Improving Decisions about Health, Wealth, and Happiness*. New Haven: Yale University Press.

Timmerman, Jens. 2004. "The Individualist Lottery: How People Count, but not Their Numbers." *Analysis* 64: 106-112.

- Ullmann-Margalit, Edna and Morgenbesser, Sydney. 1977. "Picking and Choosing." *Social Research* 44: 757–785.
- von Neumann, John and Morgenstern, Oskar. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press
- Wasserman, David. 1996. "Let Them Eat Chances: Probability and Distributive Justice." *Economics and Philosophy* 12: 29-49.