Good Reasons for Losers

Lottery Fairness and Social Risk*

Draft 15 August 2017 – Comments welcome

Kai Spiekermann (LSE)

Many goods are distributed by processes that are partly or fully random: lots are drawn, risks are distributed, gambles are taken. Once the dies are cast, there are winners and losers. The losers demand to know why they should accept the implications of being losers. In this paper, I want to focus on the losers, and especially on the reasons we can and cannot give them to justify their predicament.

First, lotteries. Oftentimes the use of randomness is a deliberate choice – we employ lotteries to allocate scarce indivisible goods. For instance, we might toss a coin to decide who among two equally deserving claimants gets the one available kidney. Many arguments have been offered for the use of lotteries, but one has been largely overlooked: Lotteries give good reasons to the losers for why they should accept being on the losing side. Namely, that a non-arbitrary procedure was used *that could have made them win*.

Second, social risk. Sometimes randomness is not deliberately chosen for fairness reasons but an unavoidable feature of distributive processes. This is so when the creation of overall expected benefits comes with expected risks. In such situations, a chancy process creates winners and losers. The losers can ask for reasons why they have been subject to losses for gains accrued by

* The idea for this paper was conceived jointly with Alex Voorhoeve. Many of the ideas presented here would not have emerged without our extensive discussions. I am very grateful for written feedback from Matthew Braham. I received helpful comments from Christian List, Keith Dowding, Bob Goodin, Alan Hajek, Geoff Sayre-McCord, Gerald Vong and many others, including audiences at Warwick University, the LSE, the ANU in Canberra, the Universities of Hamburg and Oldenburg, and a group of students from the German National Academic Foundation. All remaining errors are mine.

others. Again, we need good reasons for losers. One (but not the only) marker of such a good reason is that *the losers could have been the winners*.

Whether imposed lotteries or social risk, we end up with winners and losers determined by a (supposedly) random process. In the case of lotteries, a designer of allocative rules ought to care about finding a procedure that establishes a good reason why the loser cannot have the good. In the case of social risk, we should ask why the loser should accept a gamble that looked good *ex ante* but is a bad predicament *ex post*. If this is correct, then having a *theory that explains which reasons are good reasons to give to losers* would be useful. This paper contributes to such a theory.

The largest part of this paper is about lotteries. Lotteries are an instructive starting point because the design of the random process is fully in the hands of the allocator and the situation of the claimants is symmetrical before the lottery takes place. A series of examples in Section I demonstrates that some lottery designs, even if they involve the same probabilities, provide better reasons to the losers of the lottery than others. Section II reviews existing theories of lottery fairness. A new theory of lottery fairness based on reason-giving is presented in Section III. Once the theory is established for lotteries, the implications for social risk (Section IV) will be investigated. Section V addresses objections to some theories of lottery fairness arising from determinism, while Section VI concludes.

I. Three Lotteries

Imagine you must distribute an indivisible resource between two candidates Y and Z. Both claimants need the resource urgently, and, as far as can be established, their claims to it are equally strong. You can either give both an equal chance of the good in one of the ways described shortly, or give Y the resource outright. (Imagine you cannot give it to Z outright for whatever reason.) I also stipulate that you are known to be scrupulously impartial between Y and Z. Your impartiality implies that you prefer to allocate the resource by lottery rather than giving the resource to Y outright. Let us also suppose that this situation arises regularly, though with different claimants. Your task is to formulate an allocation mechanism that you use consistently.

In a moment, I will present different random allocation mechanisms. For convenience, I call all of them 'lotteries', even though, more technically speaking, they are best understood as extended random processes. I say 'extended' because the outcome property used for allocation is the result

of a (potentially long) causal chain that involves a genuine stochastic (random) process at some stage. This will become clearer once we consider the examples.

Let us assume throughout that two different probabilities are always equal for both candidates: the *objective probability* of winning, and the *ex ante epistemic probability* of winning (before the stochastic process takes place). It is, of course, possible that the two can diverge, for instance if, unbeknownst to everyone, a biased die is used; let us rule out such situations.

Almost everyone in the lottery literature assumes that equal objective probability, while perhaps not necessary[1], is sufficient for the fairness of a lottery. I aim to show that this view is mistaken. To show that it is mistaken I present a number of examples, serving as intuition pumps. Later in this paper I offer a hypothesis that systematizes all the intuitions these examples trigger. All my examples involve epistemically and objectively equiprobable outcomes, but I argue that these lotteries are not all on a par.

The first example is the paradigm case of a fair lottery:

> **Coin Toss**. Whenever you have to decide between two claimants, you toss a fair coin and allocate the good to Y if the coin lands heads (and to Z otherwise).

Here the random process leads directly to an outcome that determines the allocation of the good to either Y or Z. The second example uses a more personal but still random property:

> **Symmetric Birthday**. You know that one of the candidates Y and Z has a birthday earlier in the year than the other, and you believe, correctly, that both possibilities are equally likely. Whenever you need to allocate between two such claimants, you assign the good to the person born earlier in the year.[2]

This is a more extended random process. At some point in the past, the lottery determining birthdays, driven by stochastic parent choices, stochastic processes of conception, and so on, determined the birthdays of Y and Z. Prior to conception, the probabilities of Y or Z winning are equal, both objectively and epistemically. Now, many years later, the stochastically determined birthdays are used to allocate the good.

[1] Perhaps not necessary because for many theories of lottery fairness epistemic *ex ante* probability without objective probability suffices, as I will show below.
[2] For similar cases using physical characteristics, see Sher 1980.

In many ways, Symmetric Birthday is similar to Coin Toss. The initial random process has equal objective and *ex ante* epistemic probabilities. What differs is that the result is already determined long before it is used. I will argue that if both Coin Toss and Symmetric Birthday are available as allocation mechanisms, Coin Toss is preferable. This is because the loser of Symmetric Birthday is owed a reason why he should accept the result. A good reason would suggest to the loser that he *could have been the winner*. That is credibly true when using Coin Toss: the smallest perturbation of the coin would have reversed the outcome so that the loser was a whisker away from winning. The same cannot be said about Symmetric Birthday. For the outcome to be reversed, at least one of the two claimants needs to have a different birthday, and if one or both of them had a different birthday, the world would have taken a different causal path many years ago, with changes rippling through and wide-ranging effects likely. Perhaps the loser could have been the winner, but the reversal of the outcome was not a whisker away. Worse, if the loser had a different birthday, he would, short of minor shifts due to early or late birth, be the product of a different sperm and egg, so 'he' would not be the same person at all.

This problem becomes even more severe for one claimant in a variant of the Birthday case:

> **Asymmetric Birthday**. Whenever you need to decide between claimants Y and Z, you give the resource to Y if Y is born in the second half of the year (July to December), otherwise to Z. Z's birthday is irrelevant for the decision.[3]

This example does not change any of the probabilities. The only new feature of Asymmetric Birthday is that the allocation is now tied up closely with the identity of Y, but not with the identity of Z. I claim that if Y loses out in Asymmetric Birthday he has a stronger complaint than Z has if he loses out. Why? Because if Y loses, there *wasn't any genuine possibility for Y as person Y to have won*. In other words, there is no nearby possible world in which Y would have won. In fact, someone named Y and with a different birthday could have won, but that is someone quite different from the actual Y. Admittedly, it is possible to be the same person, genetically, if the birthday changes by a few days. To avoid this complication the claims made here should be read as restricted to people born deep into the second half of the year; in November, for example. Note that this is not true to the same extent for Z. Suppose Z loses out. He could have won as the actual Z in a – from his perspective – reasonably nearby possible world in which Y had a

[3] Also assume that Y and Z are not twins. If they were, the relevant asymmetry could not arise, as will become apparent later.

different birthday. This means that Y has a legitimate complaint against using Asymmetric Birthday, but Z not so much.

One set of objections needs to be addressed straight away: one might claim that there are further, hidden lotteries in the process of allocation. In particular, there might be a lottery involved in the role assignment of Y and Z, an *ersatz* lottery that works just like Coin Toss, neutralizing the problems with the Birthday lotteries. To rule out this complication, assume that the role assignment of Y and Z is not itself a lottery. Also, to keep our attention trained on the lottery that matters, it is useful to make the names Y and Z dependent on a criterion that is clearly normatively unsuitable for assigning the resource: for example, call the wealthier of the two claimants Y and the less wealthy Z.[4] Unless you as the allocator systematically want to benefit the wealthier of the two claimants, you should use a lottery to avoid benefiting wealthy Y even more.

Related problems might arise if the mechanism used for allocation changes quickly, introducing an implicit lottery about the choice of allocation process in each single case. It is best to avoid such complications by assuming, as I have above, that all decision rules are long-standing and consistently applied in the same way. That means, we assess them against the background assumption that they have been and will be used for a long time and that no further lottery is hidden in their application.

In the next section, I review the existing literature on lottery fairness before I develop my own theory. If successful, the theory will systematize the intuitions about Coin Toss, Symmetric Birthday, and Aymmetric Birthday.

II. When and Why are Lotteries Fair?

When discussing the use of lotteries in allocation problems, three questions can be asked:

1. When are lotteries a fair (or the fairest) way to allocate goods?
2. Why are lotteries fair, when they are?
3. Which lotteries are fair?

[4] And assume that Y and Z are never exactly equally wealthy.

The first question is the most familiar one from the literature on lotteries. The second question is about the functions lotteries serve, and how these functions make lotteries fair. The third question has been raised much less often. Perhaps this is because the answer is assumed to be clear: the fairest lottery is the one that gives everyone an equal chance. There is some disagreement about whether the probabilities of interest are epistemic or objective, but to my best knowledge no one has argued that any property beyond the epistemic or objective probability distribution matters to assess the fairness of lotteries. I argue that this view is mistaken. But before I do this, I quickly need to address the first two questions to prepare the ground.

When Are Lotteries Fair?

There is widespread agreement about one fundamental necessary condition for the applicability of lotteries:

> **Practical Indivisibility**. The good to be distributed cannot be divided, or any attempted division leads to side effects so bad that it ought not to be divided.

The paradigmatic cases for indivisible goods are donor organs, pharmaceutical drugs that need to be taken in threshold quantities to have any positive effect, or spaces in a lifeboat. Some other goods can be divided in theory, but are practically indivisible. In principle, instead of conducting a Vietnam draft lottery, the US government could have drafted *all* young men and deploy every single one of them for a very short time. But this would not only have been logistically difficult and costly, it would most likely have led to many more deaths and injuries.

There is also wide agreement about the second condition:

> **Scarcity**. There are more claims to the good than there are units of the good.

For example, there are fewer donor kidneys than patients with kidney failure, there are fewer volunteers than soldiers needed, there are fewer spaces in the lifeboat than passengers, and so on.

The scarcity condition mentions 'claims'. John Broome (1984, 1990) explains that these are reasons for giving the good to the different claimants. More precisely, such claims are based on duties owed to the claimants. These duty-based claims are always to be evaluated relative to the

claims of others as a matter of fairness. Fairness creates new claims, namely the fairness claim towards satisfying claims 'in proportion to their strength' (Broome 1990, 95, italics omitted). Sher (1980) and Wasserman (1996) maintain that lotteries require exactly 'equal claims' to justify the use of a lottery. The idea is that the lottery acts like a 'tie breaker' – and if the claims are not exactly equal, no tie needs to be broken. Broome disagrees, suggesting that lotteries can be the fairest allocation method even if the claims are only roughly equal.

Suppose that *need* establishes duties owed to the subjects, with corresponding claims. Now we face two patients suffering excruciating pain and we have only one dose of an effective painkiller available. If we think both patients are equally in need, everyone agrees that we should toss a fair coin. But what do we do if one of them is in a little more pain? If one thinks that lotteries can only be used as a tie-breaker, we should give the painkiller to the patient who is suffering a bit more straight away. Broome, by contrast, thinks that it would be more appropriate to have a lottery weighted towards the patient with more pain because that is the best attempt to satisfy the claims in proportion to their strength. He even suggests that, if weighted lotteries are not feasible, a lottery with equal probabilities may be the right allocation device if the consideration of additional expected harm is outweighed by the importance of respecting the claim of the patient with slightly less severe pain. To allow for Broome's view, the third condition can optionally be read in a more permissive way:

> **(Almost) Equal Fairness Claims**. All subjects have (almost) equal claims to the good(s) to be allocated.

For practical purposes, I will nevertheless assume claims of equal strength in all examples discussed in this paper.

Finally, for lotteries to be the appropriate allocative method, respecting equal claims must be more important than any other normative properties:

> **Primacy of Fairness**. The reasons for respecting claims by means of a lottery outweigh any other normative reasons supporting other allocative principles.

One important function of Primacy of Fairness is to block maximizing goodness as an alternative to lotteries. If the argument for using a lottery is to succeed, respecting duty-based claims must be more important than maximizing goodness; otherwise one should simply choose the allocation causing the greatest good, a consideration that will almost always suffice as a tie breaker.

There may also be other normative considerations that could outweigh respecting fairness. For instance, many public positions ought to be awarded according to merit. If so, it is normally required to offer the position to the best candidate. Any claims that other candidates may have towards the position are outweighed by the criterion of merit, and a proportional satisfaction of such subsidiary claims would neither be demanded nor permitted by fairness. In such cases, Primacy of Fairness does not hold.

Why Are Lotteries Fair

There are at least three different functions of lotteries: they can help to avoid partiality and arbitrariness, they can express an attitude of impartiality, and they can – though this is more controversial – have a distributive function. I go through these three functions in turn.

First, lotteries have a distinct advantage as an allocation mechanism: because they are unpredictable and, if conducted properly, impossible to manipulate, they prevent partiality by the allocator or anyone attempting to bias the decision mechanism. This exclusion of inappropriate reasons influencing the allocation has a 'sanitizing' effect (Stone 2011, section 2.4).[5] Lotteries are therefore the best decision method when any undue influence to benefit one claimant is undesirable. A lottery is also preferable over an arbitrary decision because the impartiality of the lottery is not just contingent but *guaranteed,* while arbitrary decisions may or could always be obliquely influenced by partial motives.[6]

The sanitizing effect works best if the alternative to the lottery is an unconstrained arbitrary choice by the allocator. Giving the allocator a free pass to allocate for any reason he likes invites choices for the wrong reasons. However, if the allocator is more constrained in the options available, the sanitizing effect does not work so well. Suppose, as before, that a good can either be given to claimant Y outright, or a lottery between claimants Y and Z can be used (giving the good to Z outright is not possible for whatever reason). The allocator ought to choose the lottery for fairness reasons. But the lottery does not fully sanitize: it is possible that the allocator chose the lottery not for fairness but because she is partial to Z. Therefore, the sanitizing function alone cannot explain why the lottery is fairer than giving the good to Y outright.[7]

[5] Wasserman (1996) calls this effect 'prophylactic'.
[6] A similar view is defended by Sher (1980).
[7] I am grateful to Alex Voorhoeve for bringing this problem to my attention.

A notable property of the sanitizing view is that it only demands the equality of *ex ante* epistemic probabilities. Lotteries can sanitize even if objective probabilities are not equal, and even if the *ex post* epistemic probabilities are not equal – for instance, if the subjects and the allocator discover later that the lottery was in fact not fair: the coin thrown was biased, the die was unbalanced, the lots not properly mixed, and so on. If no one was able to know of the bias *ex ante*, then no manipulation, no influence of untoward motives was possible. Even using a coin that is known to be biased can sanitize, as long as no one knows the direction of the bias before the head and tails are associated with outcomes and the coin is tossed. The sanitizing theory of lottery fairness permits the use of such purely epistemic ignorance. In the same vein, Wasserman (1996, p. 33) suggests to 'exploit impending discovery', such as using the count of stars in a newly discovered cluster as a quasi-random device. In those cases, the *ex ante* epistemic probabilities will be balanced while the objective probability is 1 for one event and 0 for all others.[8] Perhaps more puzzling for some, Wasserman's view would even sanction 'lotteries' such as giving the good to 'the person with the largest pancreas' (Kornhauser and Sager 1988, 490), as long as the *ex ante* epistemic probabilities are truly equal for everyone that could have influenced the outcome.

Second, lotteries can also have an expressive function, signalling to the possible recipients that they are treated in an impartial manner (Wasserman 1996). The use of the lottery signals to the claimants that the allocator takes their claims seriously, in a manner that respects the claims in proportion to their weight. Since my own view has an important expressive component, I will say more about this below.

The distributive function of lotteries, third, is controversial. It suggests that lotteries give something to claimants that they care about. When faced with equal claims, an indivisible good and scarcity, fairness requires that we 'go some way toward treating the candidates equally: we can give them all an equal chance' (Broome 1984, p. 45). This interpretation of the distributive function also suggests that *ex ante* epistemic equiprobability is not enough. Why? Because giving someone an equal chance requires more than just believing that one gives someone an equal chance. Broome (1984, 1990) therefore rejects the idea that the fairness of a lottery depends only on *ex ante* epistemic probabilities. While he does not explicitly discuss the possibility of epistemic and objective probability coming apart, he does say that a candidate has a complaint if her winning was 'never on the cards' (1990, p. 98). This complaint applies, for example, when using the size of the pancreas to determine the allocation decision.

[8] Though that is only true *after* the stochastic process leading to the creation of the stellar cluster has taken place.

Broome's argument for the distributive view is based on the obligation to fairly respect claims. If all candidates have roughly equal claims to a good but we cannot give that good to everyone, then fairness demands that we give 'surrogate satisfaction' (1990, p. 98) – the chance to obtain the good. This is, of course, not as good as receiving the good (after all, one cannot 'eat chances', as Wasserman 1996 points out), but it is the best one can do to satisfy the competing claims.[9]

III. You Could Have Won

In this section, I set out a theory to explain the intuitive reactions to Coin Toss, Symmetric, and Asymmetric Birthday. It is useful to define some terms to state the theory clearly. Let the *realized winner* be one of the two possible winners {Y, Z}. The *actual world* is the world after the allocation decision has been taken and the winner is fixed. The actual world can either be realized as a *Y-world* (in which Y wins) or a *Z-world* (in which Z wins). A *converse world* to an actual world is a possible world in which winner and loser are swapped.

The core claim of the theory is this. The allocator has the task to choose a lottery mechanism that does not only have a fair epistemic probability distribution; it should also provide good reasons to the loser. In particular, the allocator ought to make sure that the loser cannot reasonably reject the following argument:

> *'The lottery used was fair because it gave both of you an equal chance to the good. This means that the result could easily have gone the other way and you would have won. Moreover, I would have given the same reason to the current winner if he had been the loser.'*

To make this argument as plausible as possible, the allocator ought to choose or construe the lottery such that (for both possible losers) the actual loser could easily have been the winner. This means that if the world ends up being a Y-world it could easily have been a Z-world, and vice versa. And, for fairness reasons, the allocator ought to ensure that he treats both candidates equally with respect to that desideratum. So the reasons that apply to Y in a Z-world should be as strong as the reasons that apply to Z in a Y-world. Fairness demands the symmetry of reasons given to the loser between actual world and converse world.

[9] Kornhauser and Sager make a similar point, suggesting that a lottery's justness could be grounded on the fact that the lottery allows an 'equal division of a good that is otherwise indivisible' (p. 496) and that equiprobability creates this sort of equality in expectations.

Why is giving such a reason important? In a lottery, we allocate at least three things: (i) the expected value of the good *ex ante*; (ii) the good, (iii) the strength of the reasons for the outcome *ex post*. Virtually everybody agrees that the expected value *ex ante* should be equalized. As far as the good itself is concerned, it will have to be given to one of the claimants because it cannot be divided. Once that has been done, the expected value is zero and drops out of the picture for the loser. The only thing we can give to the loser that is still of any value *ex post* is the best possible reason for her being in this regrettable position. And indeed, given that she is badly off, being given a good reason is what she is entitled to expect.

Why is this the right reason to give? Because it captures an important aspect of fairness in allocation processes. It suggests that a win for the loser was 'on the cards'. Looking back, it means that the hope to be the winner was not based on an illusion. Looking forward, it means that the loser can live with the knowledge that the allocator chose an allocation process that nearly made him win, an important sign that his legitimate claims were taken into account. It therefore serves an important expressive function.

Note that I do not make any claims as to whether the loser of a lottery would *actually* accept the outcome when provided with the argument sketched above. The claim is more restricted: the allocator has to give a good reason – ideally the best possible reason – to the loser such that a reasonable loser ought to accept it. Since the quality of the reason differs depending on which lottery the allocator uses, the obligation to provide a good reason has implications for the fairness of the lottery chosen.

'You could have been the winner' is not the only reason we can offer the loser. Another reason is that 'you agreed (or reasonably ought to have agreed) to the lottery *ex ante*'. This is true, but does not settle the matter decisively. One response to the *ex ante* consensus argument is to point out that the sort of reason that my theory proposes adds to the overall strength of the reasons for the loser. It is more plausible to say to the loser that she agreed to the lottery in advance *and* she could have been the winner, rather than saying just the former. A second response is to deny that *ex ante* consensus is a good *ex post* reason if it is not true that the loser could have been the winner. The loser might say: 'It is true that I agreed (or reasonably ought to have agreed) to use the lottery *ex ante*. But my agreement was reasonably premised on an assumption that turned out to be incorrect: that I could have won that lottery.' If the agreement was premised that way, then the agreement is faulty because consent was given based on a false assumption, an assumption that the loser legitimately believed to be true at the time.

The reasoning set out so far suggests three desiderata for lotteries to provide good reasons to the loser:

> **Identity Preservation.** For both possible losers, it must have been plausibly possible to have a reversal of fortunes such that the *same person* who is the current loser would have been the winner.
>
> **Easy Reversal**. The outcome could have been different in a plausible, easy way.
>
> **Equal Reversal**. Easy Reversal should be equally satisfied regardless of who wins and who loses.

Identity Preservation is a necessary condition for the reason given to the loser being a good reason. If the loser could not have won the lottery as the person that she actually is, then any claims that she could have been the winner is not only cold comfort, it is simply wrong. Easy Reversal is a measure of the strength of the reason given to the loser. If the converse world would have been very similar to the actual world from the perspective of the loser then 'you could have been the winner' is plausible and thus a strong reason. The reason is less strong if the converse world differs in many respects from the actual world. Finally, Equal Reversal is a fairness criterion. Both claimants have, *ex ante*, an equal claim to the good. Therefore, they also have a strong claim to be provided with a reason (in case they lose) just as strong as the reason that would have been provided to the loser in the converse world. Using a lottery that treats the two claimants unequally in terms of the strength of the reason is unfair.

These desiderata partially order the different examples discussed above. Coin Toss satisfies all three desiderata best and is therefore the lottery that provides the best reason in the most equal way. Symmetric Birthday satisfies Equal Reversal, but does less well on Identity Preservation and certainly rules out an Easy Reversal. Having a different birthday is normally identity-changing so that plausible converse worlds do not contain the loser of the actual world. Due to non-identity effects, it may be very unlikely that the converse world contains claimants of the same identity.[10]

[10] A lot hinges on what is 'plausible', of course. But here I follow the extensive literature on the non-identity problem and assume that changes in the distant past would have changed all identities in the present with physical necessity.

It is therefore not convincing to say to the loser that '*you* could have won' because 'you' would not have existed in most plausible converse worlds.[11]

In Asymmetric Birthday, the Equal Reversal desideratum is not satisfied. This example brings out an under-appreciated fact: the difficulty of conceiving of a converse outcome can be person-relative. In this example, Y is in a special situation: his personal identity is so closely causally tied to the outcome that it is hard to see how the outcome could have been different for him. Whoever the person is that would be the winner in the converse world is someone with a different birthday and therefore not identical with Y. For Z, by contrast, nearby converse worlds are easy to imagine – the only change needed is the change of one fact: the birthday of a stranger. What is hardly possible for Y is easily possible for Z.

Asymmetric Birthday is less fair than Coin Toss and Symmetric Birthday because the Equal Reversal desideratum is violated. Y and Z will have very different *ex post* views on how things could have gone differently and the reasons that could be provided to them as losers differ in strength. But Asymmetric Birthday is good in one respect: it satisfies Identity Preservation and (to some extent) Easy Reversal for claimant Z, whose birthday is not tied to the outcome of the lottery.

The upshot is that Coin Toss is the best lottery, meeting all desiderata. Symmetric and Asymmetric Birthday are inferior to Coin Toss in different ways. Asymmetric Coin Toss is unfair because it provides reasons of different strength to the two possible losers. Symmetric Birthday is not unfair, as it treats both claimaints symmetrically. Still, it is a less good lottery than Coin Toss because it provide weaker reasons to each possible loser than other lotteries could have done.

My theory departs from a common position in the lottery literature: that the timing of the random process in the sequence of events making up the lottery does not matter. Jon Elster, for example, asserts that '[f]airness does not require the random part in the sequence of events leading to the final decision to occur as late as possible, since one random event in the chain is enough to confer randomness on the outcome' (Elster 1989, 44). The principle that Elster describes here is sometimes called the 'principle of convolution' (Kornhauser and Sager 1988, 486). It is mathematically correct, of course, that one genuine random event to make the overall

[11] There may be nearby possible worlds in which one's own birthday does not change but that of the other candidate does. This line of reasoning is stronger for individuals born mid-year. It is less strong for individuals who are born, e.g., mid-December. It is quite unlikely that the other claimant is born even later in the year.

outcome random. But randomness is not all we want from a lottery – providing good reasons to losers is important as well.

It is time to take stock of what has been achieved up to this point. The results so far provide a new normative theory of lotteries based on reason-giving. Lotteries allocate expected values of the good *ex ante*, the good, and the strength reasons that can be given *ex post* to both possible losers. Since the good cannot be shared, providing a good reason to the loser is important – it is an expression of respect for the person who had a claim but receives nothing. The allocator can facilitate strong reasons to be given to losers by making the assertion 'you could have won' as plausibly true as feasible for both possible losers. This leads to three desiderata: Identity Preservation, Easy Reversal, and Equal Reversal. These desiderata ensure that the same person could have won, that the converse world where winner and loser are swapped is nearby, and that for both possible outcomes, the converse world is equally close, so that both possible losers would get reasons of the same strength.

IV. Psychological Connectedness

The strength of a reason provided to a loser depends on whether the loser can plausibly see herself in a converse world as the winner. That is clearly difficult if the lottery used is likely to change the identity of the claimants. But structurally similar considerations apply when the converse world is one in which the same person can exist but the psychological connectedness between the loser and his counterpart in a converse world is weak.

This will be much clearer after considering an example. Until recently, some oversubscribed degree programmes at Dutch universities were allocated by lottery. For example, when applying for a degree in psychology it might be decided by lottery whether a student ends up in Amsterdam or Groningen. Let us assume that the lottery assigns a unique natural lottery number to each applicant in this process, using a random device. Now consider a proposal to use these numbers again years later, namely whenever there are two claimants for one donor kidney:

> **University Lottery Transplant**. Across the Netherlands all hospitals are advised to give the kidney to the claimant with the lowest university lottery number if there are fewer kidneys than equally deserving claimants (and they all have such lottery numbers[12]).

[12] The procedure used for claimants without such lottery numbers does not matter here.

Now consider claimants Y and Z. Suppose both are 50 years old, having graduated from university 25 years ago. Y went to Amsterdam and Z to Groningen, and both were lucky to have been assigned to their preferred cities. Y and Z find each other competing for the same kidney, and the decision is made by their old lottery number. Y wins. Z has a very high number, perhaps one of the highest number in the country, and therefore loses.

Z, I maintain, has a complaint against using the university lottery number to allocate kidneys. The problem is not with the probabilities, as Y and Z had an equal chance. The problem is, again, that the reason given to Z is not very convincing when using the university lottery number. 'You could have been the winner' rings hollow for Z. Why? Note first, since Z has a high lottery number, the result would very likely have been the same even if the other claimant had not been Y. The relevant converse worlds are therefore those worlds in which Z had a lower lottery number. But if the lottery number had been lower then, assume, Z would not have been assigned to Groningen. And if he had studied somewhere else, he would have made different friends, dated different partners, led a different live and ultimately be quite a different person from the one he is now. If that is true, then to tell Z that 'you could have won' the lottery for the kidney is far-fetched. Someone who is in some sense identical to Z, but psychologically only weakly connected to Z, would have been the winner. Z may find this quite a poor reason, and he would be right to feel that way.

The wider lesson from University Lottery is that the Easy Reversal *desideratum* is important even if what is at stake is not the identity but the psychological connectedness between the loser in the actual world and the winner in the converse world. Can it be a coincidence that all the lotteries we actually use make sure that the similarity between actual and converse world is small?

This extension to psychological connectedness concludes the discussion of lottery fairness. In the next section, we will see that the requirement to provide good reasons for losers has important applications for the normative analysis of social risk.

V. Social Risk

The recent debate on contractualist approaches to social risk imposition (Frick 2015; Kumar 2015) has structural similarities to the theory of lottery justice developed in this paper. Like a lottery, social risk involves a random process that creates winners and losers. Whether the imposition of the risk is justified depends partly on whether good reasons can be offered to the

losers of that process. In fact, providing good reasons to losers is at the very heart of the contractualist approach to social risk.

Typical situations of social risk arise when an allocator can take an action that affects many people such that each subject will accrue a benefit with high probability and a loss with low probability.[13] In the lottery case, randomness is a chosen feature of the allocation mechanism to promote justice. In the case of social risk, by contrast, randomness is an unavoidable feature of a course of action offering benefits with risk. The question contractualists are asking is: are there reasons that can be given to losers that the losers cannot reasonably reject (Scanlon 1998)?

Frick (2015) discusses vaccinations, a classic case of social risk. A large benefit of prevented deaths or illnesses for many comes with a low probability of severe harm or death for those few unfortunate with vaccination complications. Consider, just as Frick does, a hypothetical vaccination programs that affects 1 million children. All children will be killed with certainty by a virus without the vaccination. Call this course of action

No vaccination. All children die.

One vaccine is 100% reliable but also comes with a certain severe side effect. It serves as a

Baseline. All children survive, but they all have one leg permanently paralyzed.

All other vaccination programs involve social risk. The difference in the programs is about whether the identity of victims is fixed, and whether the identity is known:

Victim Lottery. All children have a 999/1000 chance to live without complications and a 1/1000 chance to die from freak complications, determined by a genuinely random process that takes place after the vaccine has been given.

Fixed Victims. 999,000 children will live without complications, but 1000 'doomed' children, whose identity is known, die with certainty, because of their genetic setup.

Frick, being an *ex ante* contractualist, wants to defend the choice of Victim Lottery over Baseline. Many *ex post* contractualists disagree. They appeal to the widely held intuition that Fixed Victims is seen impermissible when Baseline is available because it is not permissible to sacrifice 1000 identified individuals for the greater good of the 999,000 others. They then assert that Victim

[13] It is typically assumed that these individual events are statistically independent of each other.

Lottery is just as bad as Fixed Victim, on the grounds that, due to the law of large numbers, we are fairly certain that about 1000 victims will be determined by the lottery, leading to an outcome structurally similar to Fixed Victims. For *ex post* contractualists, all that matters is the existence of victims with a complaint at the end. *Ex ante* contractualists disagree – for them, it is only complaints based on information available *prior* to starting the allocation procedure that matter, which makes the choice of Victim Lottery over Baseline permissible.[14]

For this paper, however, I want to set this debate aside and instead focus on a case Frick mentions only indirectly and in passing:

Someone is Doomed. It is known that 1000 children are genetically 'doomed' to die when exposed to the vaccine, but it cannot be determined in advance who these children are because there is not enough time to perform the analysis.

Frick suggests that such cases are structurally similar to Victim Lottery because, at least for Frick, there is no difference between the purely epistemic chanciness of Someone is Doomed and the objective chanciness of Victim Lottery. Is Someone is Doomed just as good a choice as Victim Lottery? Frick suggests it is: 'Holding constant the fact of actually suffering a harm, it is typically no better for a person that she had some chance of escaping the harm rather than being certain to suffer it' (Frick 2015, 184). It is this claim I disagree with.

Frick defends his view by looking at the *ex ante* reasons that can be given to the losers of the vaccine programme. My objection to Frick's analysis is very similar to the objections against Symmetric and Asymmetric Birthday above. For the 'doomed' children, winning was 'never on the cards', in Broome's felicitous phrase. Unlike Frick's thoroughly *ex ante* approach, I maintain that reasons after the allocation has proceeded are at least as important as the reasons given before. This is easiest to see by modifying the vaccine examples again. Assume that, instead of dying, unsuccessful treatment causes the children to develop a permanent major disability and chronic excruciating and untreatable pain. This change brings out the importance of giving and receiving good reasons that are valid *ex post*, after the treatment has failed: since the children will live, the salience of giving plausible reasons shines through more clearly.[15]

[14] Assuming that having a paralyzed leg with certainty is much worse than the expected value of the gamble taken in Victim Lottery.
[15] Can it be a coincidence that *ex ante* contractualists use examples in which the losers die? Dead losers do not ask for reasons!

After the fact, the victims will ask for reasons why the risk was taken that has caused them leading blighted lives. The best answer one can give for Victim Lottery is by now familiar: '*The result could easily have gone the other way and you would have benefited and be cured. Moreover, I would have given the same reason to the current beneficiaries if they had been one of the harmed'*. Equally familiar is the observation that this statement is simply untrue when the victims were pre-determined, as is the case in Someone is Doomed. 'You' could not have been the winner – to be the winner, 'you' would have needed different genes, and that person would not have been you.[16] Frick overlooks this important difference.

In fact, Frick argues that only epistemic probabilities can ever matter in these cases:

> '… is there any greater self-interested reason to buy a ticket for a lottery where the prize will be allocated by a genuinely indeterministic mechanism [a "quantum randomizer"] than to buy a ticket for a lottery with the same epistemic odds where a ticket is either sure to win or sure to lose, in the objective sense [a scratch-card lottery with a preagreed winning number]? Again, it is hard to think of any such reason. If you had a ticket to the indeterministic lottery, and someone offered to trade it for an equivalent ticket to the deterministic lottery plus a small amount of cash, there would be no good reason to refuse this offer. …
>
> Nor could it be said that if the gamble goes badly, it is worse … to become disadvantaged as the result of a causal process that was deterministic as opposed to indeterministic. The prudential value of receiving a chance of some good, whether epistemic or objective, is strictly parasitic upon the value of the good itself.'[17]

This passage brings out my disagreement with Frick clearly. Frick asserts that it is rational to value deterministic and indeterministic lotteries in exactly the same way. But this is just an assertion, not an argument. People often have preferences not just over outcomes but over the processes leading to outcomes. It remains an empirically open question whether subjects do or do not value indeterminacy in lotteries when major losses are at stake. In addition, even if subjects are indifferent over such different lotteries as assumed by Frick, it is not clear whether satisfying these preferences is all the allocator should be concerned with. If both lotteries are equally valuable for the subject, but one lottery provides better reasons for the loser, there are

[16] For reasons that are easy to anticipate I also disagree with Frick's discussion of 'The Argument from Determinism' (197-201). His remarks about lotteries and scratch cards are diametrically opposed to my view set out in the next section. However, for space reasons I do not pursue this disagreement here.
[17] The square brackets are Frick's, but the ellipses are mine.

good grounds to choose the latter lottery, regardless of whether the individuals are *ex ante* indifferent betweem them.

There is a wider lesson to be taken from this. The distinction between *ex ante* and *ex post* contractualists is too coarse. It misses the distinction between the time of reason-giving and the time at which the facts these reasons appeal to are true. My view is in favour of *ex post* reason giving, but these reasons appeal primarily to facts that are made true *ex ante*, especially modal facts about the lottery. Note well that I do not endorse the standard *ex post* view that gives losers a veto merely based on the fact that they are the losers. This version of *ex ante* contractualism, Frick is right to point out, would lead to highly revisionary and quite unattractively status quo biased prescriptions. But I do argue that giving good *ex post* reasons is important.

In section III, I argued that we allocate three things: (i) the expected value of the good *ex ante*; (ii) the good, (iii) the strength of the reasons for the outcome *ex post*. The first two points become irrelevant for the losers after the allocation has finished. What remains *ex post* is nothing but the reasons the allocator can give them for their predicament. It is these *ex post* reasons that are the sign of the victims of social risk being treated with as much respect as possible. Since society has nothing else to give, providing the best possible reasons *ex post* is of moral importance.

VI. The Challenge from Determinism

Can we make any sense of the idea of distributing chances if we assume determinism? If a system is deterministic this means that the future trajectory through the state space of the system is fixed by the current state (and possibly also the past states) of the system. Determinism about the whole world means that the future of the world is fixed, given the history of the world. All possible future events are then already determined: they will happen or not. Consequently, the objective probabilities over possible outcomes of all future lotteries are either 0 or 1 – you either win or lose with certainty (e.g., Sher 1980, 204; Otsuka, n.d.).

We don't know for sure whether determinism[18] is true or not. Our best physical theories suggest that determinism about the macro-physical world is likely true at least locally and at least most of the time, but that there might be genuine randomness in the microphysical world. In any case, even if the physical world were fully deterministic on some level, it is far from clear whether such

[18] I omit the addition 'about the world' from now on.

a physical determinism would make the notion of objective probability obsolete (e.g., Hoefer 2007; Ismael 2009; Briggs 2010). One way to defend objective chance is to make it level-specific and to show that on the level of human agency determinism is false and objective chance exists, while on a lower level of description determinism may be true and objective chance does not exist (Glynn 2010; Emery 2015; and especially List and Pivato 2015 for a concise review of the debate and a formally developed proposal).

If determinism is true then a Laplacian demon will be able to predict the outcome of every lottery with certainty. From the perspective of the demon all objective probabilities are 0 or 1. But for a human observer, even for a human observer with the best possible evidence a human observer can possibly gather, a well-constructed lottery is not predictable. Take a coin toss, for example. Assume that the coin is thrown high enough, with a strong initial spin. Also assume that the coin is allowed to hit the floor and bounce until it has settled on one side.[19] To predict the outcome, the observer needs to have the starting position of the whole system with infinite precision, including the position of the coin, data about the initial forces from the throw, the atomic structure of the coin down to every single atom, any air movements, the position of all molecules, the surface on which the coin lands, and so on, and so forth. Any smallest deviation could lead to a different result. Good coin tosses are effectively unpredictable. (Diaconis, Holmes and Montgomery 2007)

This unpredictability is a hallmark of chaotic, non-linear systems. Importantly, the unpredictability does not depend on whether the system is indeterministic or not. Research into chaotic systems has provided us with many examples of deterministic systems that are unpredictable because of tiny changes in the initial conditions (e.g. Smith 2007). Prediction is impossible for three reasons. First, it is impossible to measure all initial conditions with sufficient precision. Second, since real-world systems are normally not completely closed, we would have to measure all outside influences. Third, even if we could perform all the measurements, we neither have the precise mathematical models nor the computational power to perform the required calculations to determine the trajectory of the system. Any tiny deviation in initial position, model specification or even the tiniest error in calculation would throw us off course.

This shows that there is an important difference in the epistemic accessibility of the results of a coin toss that is about to happen, and, for example, a scratch card lottery.[20] According to our

[19] This is surprisingly important. Without these additional random shocks there is an element of predictability left, as Diaconis, Holmes and Montgomery 2007 show.
[20] I owe the scratch card example to Mike Otsuka (n.d.). It also features in Frick (2015), as seen above.

best scientific knowledge of the phenomenon, no human being will ever be able to predict a well-executed coin toss, even with the best measurement devices. I take this to be a nomological necessity, based on our best theories of the natural laws that govern chaotic, non-linear systems. By contrast, the scratch card can be predicted – either, trivially, by observing it prior to coating, or with a somewhat sophisticated scanning device to peek through the coating.

A paradigmatic case of an epistemic probability is about an outcome that, while not known by the relevant subject, could be known by uncovering the relevant information, as in the scratch card case. What about those cases where the outcome is *ex ante* not knowable for any human observer, given the laws of nature? I will call the probabilities attached to such outcomes objective probabilities, to contrast them with the paradigmatic cases of merely epistemic probabilities. The difference between objective and epistemic probability is then about the knowability of the actual outcome given the constraints created by chaotic, non-linear processes.

For example, for me to throw a six with a fair die (and a proper, strong and therefore unpredictable throw) has an *objective* probability of 1/6 in this sense. But if someone else throws the same die and covers it with a cup, then, after the throw, I believe with a *subjective* probability of 1/6 that a six is under the cup. And this probability is subjective regardless of whether I can access the information. The fact that it is accessible in principle to a human observer suffices.

Naturally, some determinists will protest and point out that all probabilities are ultimately epistemic because all outcomes are knowable (namely, by a Laplacian demon). I would urge those readers to replace instances of 'objective probability' in this paper with '*ex ante* epistemic probability of an event that is unpredictable (for nomological reasons)'. This underlines the distinction between epistemic uncertainty where the outcome is, at least in principle, knowable *ex ante*, and stochastic events where that is not the case, similar to the difference between Someone is Doomed and Victim Lottery in the last section.

In fact, nothing much hinges on this terminology. What this paper does hinge on is the possible truth of statements like 'X could have won the lottery' after X lost the lottery. Perhaps some hard-core incompatibilist determinists want to deny even this. Perhaps they will say that in the world in which X loses the lottery it is simply false that it could have been any other way because there is only one unique trajectory through state space. But surely winning the lottery was at least conceptually possible for X. There surely must be a difference between X not winning in the lottery that just took place, and X not winning because there was no lottery. This is all I need.

VII. Conclusion

This paper suggests that the normative literature on lotteries is incomplete and that normative theories of lottery fairness need to be revised. Lotteries can have the same epistemic and even objective probabilities but – surprisingly – differ in the strength of the reasons the offer to the losers of the lottery. We should, *pro tanto*, prefer lotteries that provide stronger reasons to the losers, and – as a matter of fairness – reasons of equal strength to all possible losers. Being able to tell the loser that 'you could have been the winner' is an important expression of respect. The allocator should aim for a lottery that makes this statement true, and plausibly so.

What applies to lotteries also applies, I have shown, for cases of social risk. Fairness requires that we offer all possible losers reasons of equal strength, and it is better if the reasons are stronger rather than weaker. That the losers would have been willing to take the risk *ex ante* is, by itself, a weak reason if they never had a chance to benefit. Proponents of *ex ante* contractualist approaches to social risk may worry about the *ex post* nature of the reasons my theory demands. But while reasons for losers apply *ex post*, they appeal to facts that are made true *ex ante*. This suggests that the *ex ante*/*ex post* distinction is too coarse to answer the question as to what kinds of reason we owe to the victims of social risk.

References

Briggs, Rachael. 2010. 'The Metaphysics of Chance.' *Philosophy Compass* 5: 938–52.

Broome, John. 1984. 'Selecting People Randomly.' *Ethics* 95(1): 38–55.

———. 1990. 'Fairness.' *Proceedings of the Aristotelian Society* 91: 87–101.

Diaconis, Persi, Susan Holmes, and Richard Montgomery. 2007. 'Dynamical Bias in the Coin Toss.' SIAM Review 49 (2). *Society for Industrial and Applied Mathematics*: 211–35.

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

Emery, N. 2015. 'Chance, Possibility, and Explanation.' *The British Journal for the Philosophy of Science* 66 (1): 95–120.

Frick, Johann. 2015. 'Contractualism and Social Risk.' *Philosophy and Public Affairs* 43 (3): 175–223.

Glynn, L. 2010. 'Deterministic Chance.' *The British Journal for the Philosophy of Science* 61 (1).

Hoefer, Carl. 2007. 'The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance.' *Mind* 116 (July): 549–96.

Ismael, J T. 2009. 'Probability in Deterministic Physics.' *The Journal of Philosophy* 106 (2): 89–108.

Kornhauser, L. A., and L. G. Sager. 1988. 'Just Lotteries.' *Social Science Information* 27(4): 483–516.

Kumar, Rahul. 2015. 'Risking and Wronging.' *Philosophy & Public Affairs* 43 (1): 27-51.

Otsuka, Michael. N. d. 'Determinism and the Value of Fairness and of Equal Chances'.

Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge (MA) and London: Belknap.

Sher, George. 1980. 'What Makes a Lottery Fair?' *Noûs* 14(2): 203–16.

Smith, Leonard. 2007. *Chaos: a Very Short Introduction*. Oxford: Oxford University Press.

———. 2011. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford and New York: Oxford University Press.

Wasserman, David. 1996. 'Let Them Eat Chances: Probability and Distributive Justice.' *Economics and Philosophy* 12: 29–49.