Political Concepts
Committee on Concepts and Methods
Working Paper Series

16

September 2007

Conceptualizing Lotteries

Peter Stone

Stanford University

(peter.stone@stanford.edu)

C&M
The Committee on Concepts and Methods
www.concepts-methods.org

IPSA
International Political Science Association
www.ipsa.ca

CIDE
Teaching and Research in the Social Sciences
www.cide.edu

The C&M working paper series are published by the Committee on Concepts and Methods (C&M), the Research Committee No. 1 of the International Political Science Association (IPSA), hosted at CIDE in Mexico City. C&M working papers are meant to share work in progress in a timely way before formal publication. Authors bear full responsibility for the content of their contributions.

The Committee on Concepts and Methods (C&M) promotes conceptual and methodological discussion in political science. It provides a forum of debate between methodological schools who otherwise tend to conduct their deliberations on separate tables. It publishes two series of working papers: "Political Concepts" and "Political Methodology."

Political Concepts contains work of excellence on political concepts and political language. It seeks to include innovative contributions to concept analysis, language usage, concept operationalization, and measurement.

Political Methodology contains work of excellence on methods and methodology in the study of politics. It invites innovative work on fundamental questions of research design, the construction and evaluation of empirical evidence, theory building and theory testing. The series welcomes, and hopes to foster, contributions that cut across conventional methodological divides, as between quantitative and qualitative methods, or between interpretative and observational approaches.

Submissions. All papers are subject to review by either a member of the Editorial Board or an external reviewer. Only English-language papers can be admitted. Authors interested in including their work in the C&M Series may seek initial endorsement by one editorial board member. Alternatively, they may send their paper to workingpapers@concepts-methods.org.

The C&M webpage offers full access to past working papers. It also permits readers to comment on the papers.

www.concepts-methods.org

Conceptualizing Lotteries

Peter Stone
Political Science Department
Stanford University
Stanford, CA 94305-6044
(650) 725-2916
peter.stone@stanford.edu

Abstract: A number of philosophers and social scientists have investigated the use of *lotteries* to make various kinds of decisions. These investigations have not, however, produced a rigorous definition of a lottery. This paper offers such a definition, and explains how a lottery fitting this definition can be fitted into the decision-making process. It does not explore the question of why one might wish to use a lottery in decision-making, but it does point out several ways in which the definition offered here should facilitate efforts to answer this question.

An earlier version of this paper was presented at the 2006 Annual Meetings of the American Political Science Association.

1. Introduction: It is commonly accepted that there exist many circumstances under which decisions ought to be made by tossing a coin, rolling a die, or using some other form of *lottery*. Lotteries have been used, or proposed for use, in deciding what restaurant to visit for dinner; to decide whose taxes get audited; to select patients who are to receive organ transplants; and to fill jury boxes. Not surprisingly, there have arisen a number of theoretical works that have subjected lotteries to critical scrutiny (e.g., Sher, 1980; Kornhauser and Sager, 1988; Elster, 1989; Duxbury, 1999; Goodwin, 2005). This literature has explored many of the reasons why people might wish to employ lotteries to make a wide variety of decisions. But typically, these works take for granted what a lottery is, and what it means to use one in making a decision.

This paper takes a step backwards, by offering a formal definition of a lottery and by showing how lotteries get incorporated into decision-making. It sets aside the question of why one might wish to use a lottery. It regards as pointless any effort to specify reasons for using a lottery until it is clear what one is and what it would mean to use one. It suggests that clear answers to these questions will make possible a deeper understanding of the reasons why one might wish to use a lottery at various times. While our everyday understanding of lotteries can make possible very fruitful research about them, as the existing literature suggests, this paper aims at something more.

Section 2 of the paper develops and defends a definition of a lottery, a definition consisting of three parts. Section 3 elaborates upon and explores each of these parts. Section 4 uses this definition to subdivide the class of lotteries into three different subclasses. Section 5 spells out just what takes place when a decision-maker uses a lottery to make a decision. Section 6 argues that, in light of the place lotteries fill in decision-making processes, all three parts of the definition are essential to it. Section 7 concludes by raising some theoretical problems that require further attention given the definition of a lottery constructed and defended here.

2. Defining a Lottery: Throughout this paper, I shall take for granted that the purpose of a lottery is to assist in making some sort of decision. A lottery might in fact be used for some other purpose, but that is incidental to it. A hammer might be used as a paperweight, but the function served by a hammer is to strike things. Similarly, one can toss a coin purely as a form of idle amusement, but the function served by coin tosses and other lotteries is to assist in decision-making.[1]

Assume, then, a decision-maker (individual or collective) who seeks to employ a lottery to make some form of decision. Common usage suggests that any of the following processes might prove acceptable to her:

Drawing balls from an urn.

Tossing a coin.

Picking names out of a hat.

Drawing straws.

Rolling dice.

Cases such as these constitute *paradigmatic cases* of lotteries, and in offering my definition I shall treat them as such.

Any sensible definition of a lottery, I assume, must be compatible with these paradigmatic cases. It makes sense, therefore, to start with these cases and ask what properties they share. Three are readily apparent. All of these lotteries are *processes*.

[1] Here I follow Jeremy Waldron's suggestion that "Justificatory argument in political theory and jurisprudence must precede conceptual analysis, not the other way round" (Waldron, 1989, p. 69).

Each one is capable of generating a number of different *outcomes*. And the outcomes of each one cannot be predicted with *certainty*. This suggests the following definition of a lottery:

(D) A lottery is *a process capable of generating a specific set of outcomes, in which the particular outcome to be expected whenever the process occurs cannot reasonably be predicted with certainty*.

This definition includes a number of processes that may not intuitively sound like lotteries. Suppose that a volcano erupted periodically, and that each time it erupted, a given quantity of ash was released. The volcanic eruption would indeed be a process that generated one or more outcomes—in this case, quantities of ash. And the precise quantity of ash to be expected from an eruption could not be predicted in advance with certainty. This process would therefore satisfy all three criteria of my definition, and so thereby count as a lottery. This poses no problems for my definition; it is possible to imagine using a volcanic eruption just as one uses a coin toss or drawn straws to make a decision. Granted, the process could not be generated at will the way a coin toss can be, and even if it could, the costs of employing it and measuring the results would surely be prohibitive. But all that this proves is that it would be inadvisable to use such a lottery to make real-world decisions. An inconvenient lottery is still a lottery, just as a poorly designed hammer is still a hammer.

I suggest, then, that every lottery shares these three properties. In the section that follows, I shall examine each of them in more detail so as to make plain what each entails. But first, two quick points of clarification are necessary.

First, it might appear at first glance that some idea of *equiprobability* or *randomness* should be incorporated into the definition I offer. Most people, after all, would say when a coin is tossed, the chance of getting heads is equal to the chance of getting tails. But notice that this is not true of all coin tosses. Some coins, after all, are biased. But such bias should not exclude such coin tosses from the class of lotteries. A loaded die, or a coin that favors heads, is still a lottery; it is simply not a lottery that one would want to use for the same purposes as a balanced die, or a fair coin.[2] This suggests that there may be important distinctions to be made within the class of lotteries. I shall take up this subject in section 4.

Second, I shall use the term "lottery" largely without observing the type-token distinction. That is, when I speak of a lottery such as a coin toss, I shall be referring to both the general process of tossing a coin and the specific occurrence of the toss of a particular coin at a particular time and place. I shall make no general effort to avoid this ambiguity; any attempt to do so would require much circumlocution for minimal gain. Instead, I shall provide clarification only when the type-token distinction is both important to the argument and yet with the potential to cause confusion.

3. The Three Properties of a Lottery: To begin this review of the three essential properties I attribute to a lottery, I note that a lottery is a *process*, not a *device* or *object*. It is not a coin that constitutes a lottery, but the process that takes place whenever the coin is tossed. There are many things that can be done to or with a coin; all of them constitute processes, but only some of them constitute lotteries, and different processes involving

---

[2] Note that if the loaded die was so unbalanced as to *never* land on any side but one, then the outcome of tossing the die could be predicted with certainty. In that case, the die toss would not, on my definition, count. But this, I believe, fully squares with our intuitive understanding of the concept. I shall have more to say on this topic in the following section.

the same coin may count as different lotteries. One can melt a coin in a furnace; this process has one determinate outcome (a puddle of metal), and does not constitute a lottery. One can drop it downwards from a height of a few inches, or toss it upwards into the air to a height of several feet. These two processes might both count as lotteries, but not necessarily as the same lottery, despite the fact that both employ the same device. (Whether or not they should in fact count as the same lottery is an epistemic question, as will be discussed later.) It is therefore wrong to treat statements about coin tosses and the like as statements about devices (e.g., Levi, 1986, p. 48). Of course, it is often possible to speak of a device without any lack of clarity as to how the device is to be used. In such cases, it might be perfectly intelligible to treat process and device as identical. But even in such cases, the device merits being called a "lottery" strictly by courtesy.

It is also inaccurate to identify a lottery with its potential use in a given situation. Consider, for example, the following passage from a paper by George Sher: "It is generally agreed that when two or more people have equal claims to a good that cannot be divided among them, the morally preferable way of allocating that good is through a tie-breaking device, or *lottery*, which is fair" (Sher's emphasis; Sher, 1980, p. 203). Here Sher defines a lottery in terms of its function; any process[3] capable of breaking a tie between claimants counts as a lottery. While this definition clearly includes all the paradigmatic cases of lotteries—coin tosses, balls draw from an urn, etc.—it also includes much else, including many processes that intuitively do not sound like lotteries. One could break such a tie by auctioning off the good to the claimants, for example, but

[3] Technically, Sher refers to the *device* used to break a tie. The result is that he runs together in his definition of a lottery a process, the device that generates it, and the use to which it is put. It is clear, however, that Sher intends to place the emphasis upon the *use* of the lottery in his definition, and so this is where I direct my attention.

there is no reason to count an auction as a lottery. What Sher probably means is that a lottery of some sort is the *right* way to break such a tie, but this requires more of a distinction between a lottery and its use than Sher's definition makes possible.[4]

As mentioned before, a lottery is best regarded as a sort of tool. It makes sense to define a tool in terms of its use. A hammer, thus, is a tool constructed for forcing nails into wood. But there are many ways to accomplish the purpose for which a tool is constructed, and it would be a stretch to treat all of them as equivalent. One can use many things to hammer a nail into a piece of wood—a rock, a table, your head, etc.—but this does not make these objects into hammers. To qualify as a hammer is not simply to perform a task, but to perform that task *well*, or at least well enough. This makes the definition normative in a sense; bad hammers just don't qualify as hammers. But it also means that one can distinguish between the tool and its use, so that it still qualifies as a tool when it is not being so used. A hammer used as a paperweight is still a hammer, and a coin tossed for entertainment is still a lottery, even if it is not being used for decision-making purposes at the time. Moreover, one can speak of the properties of a tool that make it suitable for its intended use , and those properties will be independent of the use itself. Lotteries may be useful in decision-making, but if this is the case, it is because of some property or properties they possess, and those properties can be specified without explicit mention of decision-making *per se*.

[4] Other theoretical treatments of lotteries similarly conflate definition and use. Thomas Gataker, an English theologian who wrote the first known book on lotteries in 1619, defined a lottery as an "event merely casual purposely applied to the deciding of some doubt' (Quoted in Rescher, 1959-1960, p. 156). In a more recent treatment of the topic, Lewis Kornhauser and Lawrence Sager write that "A lottery allocates a benefit (sometimes called a 'prize') among a designated group of potential beneficiaries ('candidates' who comprise a 'pool') according to a stipulated procedure (the 'payoff condition')" (Kornhauser and Sager, 1988, p. 485). While it is true that the allocation of benefits and the deciding of doubts are important *uses* of lotteries, this fact must not be taken to yield a definition of a lottery in any straightforward way.

Lotteries thus do not exist in nature any more than hammers do. They are inherently artificial, and reflect some level of human choice and purpose. A rock might conceivably be described as a hammer, but not in a world in which no person wanted nails to enter wood. Similarly, processes might generate outcomes in a world in which people never needed them for decision-making, but so long as this was the case, it would make no sense to describe them as lotteries. This will become even more evident in considering the second condition that lotteries must possess.

Note that one and the same process may be described in various different ways. One can describe the set of possible outcomes that could result from rolling an ordinary die as "1, 2, 3, 4, 5, or 6", or as "even or odd," or as "6 or some other number," or many other ways. Similarly, one can describe the set of possible outcomes from a volcanic eruption that produces ash as "less than or equal to 1 inch deep or greater than 1 inch deep," or "at least x, but less than $x+1$, inches of ash, where x ranges over all positive integers,"[5] or in an infinite number of other ways. Under a given description of the outcomes, a process might be useful for certain purposes, but under a different description, it might not. This is easy to recognize when one takes into account the fact that a die roll may be described as having exactly one outcome—the generation of some number. But with only one outcome to be expected, the process does not qualify as a lottery at all. (One can predict with perfect confidence that every time a die is rolled on a flat surface *some* number will come up.) Thus, it is necessary to say that a given process counts as a lottery only under a certain description of the outcomes, and that different

[5] In this case, of course, the lottery has infinitely many possible outcomes. It may seem counterintuitive to describe such a process as a lottery at all, but I think it more appropriate to describe it as a lottery that is far less useful to human purposes than other processes (or even the same process under a different description).

descriptions of the same process (such that the set of possible outcomes is defined differently in each) count as different lotteries.

The set of outcomes associated with a lottery is therefore not a feature of the world independent of human judgment. Some agent must assign a set of outcomes to a lottery before it can be uniquely identified as such.[6] This assignment is partly determined by the features of the process known to the agent. If an agent wishes to toss an ordinary six-sided die, then he cannot describe the set of possible outcomes as "1, 2, 3, 4, or 5." He can generate a set of five outcomes from the die toss if he wishes, but only by treating, for example, a roll of six and a roll of some other number as a single outcome. Subject to this constraint, however, the set of outcomes can be specified any way the agent likes. This is not to say that the agent can make the specification on a whim, but rather that which specification should be selected will depend upon the purpose for which the agent in question needs a lottery. If an actor wishes to select one of two possible outcomes, then she would be well-advised to employ a lottery with exactly two outcomes associated with it.

What this discussion reveals is how thoroughly human judgment is embedded in every stage of the process of defining a lottery. Lotteries are thus not objects that exist in the world independent of human knowledge and purpose, any more than traffic laws or mathematical theorems are. Treating them as such can only lead to theoretical confusion.

Lotteries are thus a strict subset of the set of all possible processes. This subset is marked out by the fact that the processes within it are all capable of generating some specified set of outcomes. But this condition by itself is not enough to distinguish

[6] Thus, even "natural" lotteries—lotteries not subject to generation via deliberate human action, like a volcanic eruption—have an artificial component to them.

lotteries from non-lotteries. Consider the process of turning on a television. One could specify the set of possible outcomes that could result from this process as the set of channels that could appear on the set after it is turned on. But there are clearly circumstances under which it would be odd to describe this process as a lottery. If the agent employing the process knew the channel to which the television was set before it was last turned off, then that agent can perfectly predict the outcome of the process of turning it back on. It would be foolish to call such a process a lottery. However, if the agent was unaware of this information, then certain prediction of the outcome would be impossible. Under such circumstances, the process would count as a lottery as assuredly as a die roll or coin toss. This is the point established by the third condition a lottery must satisfy.[7]

This characterization of a lottery is inherently an epistemic one. It refers to what an agent can reasonably predict about the outcome of a process. Processes whose outcomes can be predicted with certainty are clearly not lotteries. Those that cannot be so predicted are. Life, of course, is never really certain; our expectations are sometimes foiled, and even processes that we believed inevitably lead to a single determinate outcome occasionally surprise us by producing something completely different. Nevertheless, even where agents lack logical certainty that an outcome is impossible,

[7] Note here that if the channel to which the television set was attuned when the set was last turned off is known, the process of turning it back on can be described in two ways. One could say that the process in question—turning the set on—is capable of generating multiple outcomes, but that due to certain circumstances, the next occurrence of the process can be predicted with certainty. Alternatively, one could describe the process as turning the set on *when the set was known to have been set to such-and-such channel*. In this case, the outcome of the process is also predictable in advance, but the process has one and only one possible outcome. The former, but not the latter, description of the process stresses its similarity to a lottery. That is, on the first description, it fails condition three of the definition for a lottery, but on the second description, it fails both conditions two and three.

they can and do enjoy a type of certainty—call it "moral certainty"—in many cases.[8] Defining such certainty is admittedly a challenge, but to ignore it entirely is to ignore a real fact about the nature of human knowledge. There are times when we take for granted—and are justified in taking for granted—that one and only one outcome is possible, and at such times it makes no sense to speak of a lottery taking place. The fact that what we take for granted sometimes fails to occur does not change this.

Note further, however, that the lack of certainty required by a lottery is not a purely subjective, psychological phenomenon. The definition of the lottery offered here requires that this lack of certainty be *reasonable*. This reasonableness is a function of the evidence obtainable[9] by the agent. If the agent *should* be able to predict with certainty the outcome of the process when it occurs, then the process is not a lottery for that agent. If the agent mistakenly believes that the outcome cannot be predicted, then the agent's identification of the process of a lottery is also mistaken. Of course, proving a negative is difficult; it is next-to-impossible to *prove* that, given certain available information, there exists no way to know with certainty the outcome of a process. But all that this establishes is that human assessment of social processes—including lotteries—is fallible.

Finally, note that the specification of the set of outcomes associated with a lottery is also governed by normative considerations. An agent may predict with certainty that a

[8] This is the right way to think about turning on a television whose channel setting is known. The television might be broken, after all. But when moral certainty exists, then the violation of that certainty can be interpreted several ways. The most common is to assume that some as-yet unspecified *ceteris paribus* condition necessary for that certainty to be justified was violated. True moral certainty typically involves something much deeper than a surface regularity—something like a physical law. Morally certain individuals are thus loathe to admit that real violations of such a law are possible.

[9] I say "obtainable," rather than "obtained," because an agent may fail to be reasonable in at least two different ways as far as information is concerned. He may fail to derive the correct beliefs given the information he has. Or he may fail to collect a sufficient amount of information, given the information already possessed and the importance of the purpose at hand. (Technically, he might also collect too much information, but such a possibility is irrelevant here.) Either fault may lead an agent unreasonably to believe that the outcome of a given process cannot be predicted. See Elster (1989, Ch I) for a discussion of the relationship between information and rationality.

coin toss will result in either heads or tails; that agent thereby excludes the possibility that the coin may land on its edge, or that it will simply hover in midair without returning to the ground. While the latter exclusion seems pretty safe, the former may not be, especially if the surface onto which the coin is tossed is very rough. Nonetheless, the agent may be certain that the excluded outcomes cannot happen in the ordinary course of things; if one of these excluded outcomes does happen, that agent will have to rethink how the world works in some sense. [10] Moreover, the certainty can be either reasonable or unreasonable; if the surface really was very rough, the agent would be acting irrationally in refusing to admit that the coin might fail to land flat on one side. And if the agent had good reason to believe that only one outcome (given the description of possible outcomes used by the agent) was possible for a given process, then the agent would be irrational in treating that process as a lottery at all.

4. Three Types of Lotteries: A lottery, on my account, should thus be contrasted with processes whose outcomes could reasonably be known in advance. In other words, lotteries are distinguished by their lack of certainty. A lottery by definition is a lottery with non-certainty. (I say "non-certainty," rather than "uncertainty," for reasons that will become clear shortly.) This suggests an obvious link between lotteries and probability. Because the outcomes of a lottery are uncertain, it intuitively makes sense to speak of the various outcomes of a lottery as happening with various probabilities. However intuitive

[10] As noted before, if the outcome of a process can be predicted with certainty, it can be described either as

1) a process capable of generating several possible outcomes, but whose occurrence under certain circumstances makes possible the certain prediction of the outcome; or as
2) a process capable of generating one and only one possible outcome (because the description of the process incorporates the circumstances that facilitate prediction).

For this reason, it makes sense that certainty would be involved in the process under either description. The only difference the description makes is to shift the locus of certainty from the characterization of the process to the characterization of the circumstances under which the process occurs upon a given occasion.

this step may be, it is important to make it carefully. There are two potential pitfalls that one must avoid in order to associate lotteries with probability.

The first potential pitfall concerns the possibility of making probability assignments. To make the link between lotteries and probability a tight link, one must believe that it is possible to assign probabilities to all outcomes associated with all well-specified processes with non-certain outcomes. It is not obvious that this can always be done. Suppose, for example, that one knows nothing about a lottery except that it is capable of generating a certain set of outcomes. Can one assign probabilities to each of these outcomes? Many probability theorists say yes, and appeal to the principle of insufficient reason. This principle states that when nothing is known that would warrant predicting which of several outcomes will be generated by a process, all of these outcomes are to be assigned equal probabilities. But the principle itself is controversial, and leads to some highly unintuitive conclusions (cf. Elster, 1989, pp. 43-44). Suppose, for example, that a die is rolled, and nothing is known about its properties. If the outcomes of the die roll are labeled “1, 2, 3, 4, 5, or 6”, thereby generating a lottery, the principle of insufficient reason implies that the probability of rolling a “1” is 1/6. If, however, the outcomes of the die roll were labeled “one, or some other number,” thereby generating a different lottery, that same principle would insist that the probability of rolling a “1” is ½. There is no reason, however, to expect that the probability of rolling a “1” would be different in the two lotteries. Given such anomalies, the only way to salvage the principle of insufficient reason would be to specify some reason for regarding some specification of outcomes as “more natural” than others. But probability theorists have not found a convincing way in which to do this.

If the principle of insufficient reason is rejected, then the only alternative is to admit that for certain lotteries—in particular, lotteries about which little or nothing is known—there is no rational basis for assigning probabilities to the various outcomes. The result is that one must distinguish between lotteries for which probabilities can be assigned to outcomes and lotteries for which such assignments cannot be made. The economist Frank Knight once distinguished between *risk*, which people face when they are capable of estimating the probabilities of various events, and *uncertainty*, in which there is no basis for such estimation (Knight, 2002). Following Knight, we can distinguish between *lotteries with risk* and *lotteries with uncertainty*,[11] depending upon whether or not probabilities can be attached to the various outcomes. (I reserve the term "non-certainty" to describe both types of lotteries.) Any probability theorist committed to the proposition that probability assignments can always be made will believe that the latter category is empty. I shall retain the category here, however, for purposes of completeness.

To act upon the assumption that probability assignments can always be made to the outcomes of a lottery is thus to presuppose a controversial position with regard to probability. This leads to the second potential pitfall in describing the relationship between probability and lotteries. The definition of probability remains highly controversial. While the dominant interpretation of probability today is the Bayesian one, this position is by no means universally affirmed. This makes a difference if one wishes to speak of a lottery generating such-and-such outcomes with such-and-such probability.

[11] I hasten to add that even with a lottery with uncertainty, it is not the case that absolutely *nothing* whatsoever is known about it. The complete list of outcomes the process is capable of generating must be known and specified clearly, for example; otherwise, the process does not qualify as a lottery at all. And this is true of lotteries more generally; they involve a complex combination of knowledge (of the process' possible outcomes) with ignorance (of the precise outcome to result from the process).

Such a claim might be defensible on one definition of probability but not on another. Defending any particular claim, moreover, would require defending a particular definition of probability.

I shall not attempt to defend any specific conception of probability here. Instead, I shall confine myself to making two observations. First, if a given process generates two or more outcomes with nonzero probability, then the outcome of that process cannot be predicted with certainty. This claim is true under any definition of probability one might embrace. It is true even under a definition that does not allow for the possibility of uncertainty, one that treats all lotteries as involving risk of some sort. Thus, any conception of probability is compatible with the definition of a lottery given above. They will all recognize any given lottery as a lottery, even if they disagree about the properties they would assign to it. Second, if two agents disagree as to the proper way to understand probability, they may disagree as to the probability of each outcome that a particular lottery might generate. This is significant if one wishes to make an argument in favor of using a lottery in decision-making based upon the properties it possesses. One might, for example, wish to argue that a tie between two options ought to be broken by a lottery in which both outcomes occur with equal probability (such as a fair coin toss). But if I have one understanding of what probability is, I may regard a given coin toss as fair, whereas you, with a different understanding of probability, may not. All that this proves is that agreement as to how probability is to be understood is a prerequisite for further argument about whether and when a given lottery might justifiably be used. I assume such agreement exists throughout the rest of the paper for the sake of argument.

This is probably an appropriate point at which to revisit the question of how to distinguish one lottery from another. As noted before, the distinction does not track the distinction between the physical objects employed to generate lotteries. Two different coins might be tossed, and each might count as an instance of the same process. Alternatively, the same coin might be tossed two different ways, and thereby count as instances of two separate processes. What distinguishes the two processes is the level and kind of knowledge people have about them. A coin tossed under one set of circumstances[12] might constitute a process about which much is known, and for which probability assignments are easy. That same coin tossed under another set of circumstances might constitute a process about which virtually nothing is known, and for which probability assignments are impossible. The first is a lottery with risk, the second a lottery with uncertainty. All that varies between the two is the level and kind of knowledge about the two processes.[13]

This implies that two lotteries are identical if and only if all their features relevant for predicting their outcomes are identical. Once again, this is an epistemic distinction. There will no doubt be unknown features that distinguish between coin tosses—even tosses of the same coin—in such a way that, were they known, it would be impossible to treat them as instances of the same process.[14] But this makes no difference so long as

[12] Ayer (1963) distinguishes between an event whose outcome is unknown (such as a coin toss) and the circumstances under which the event takes place (e.g., whether the coin toss takes place on the moon or on earth). On the account offered here, this distinction does not exist. A coin toss on the moon constitutes a different process than a coin toss on earth if different probability assignments can be made to each of them.

[13] Indeed, a change in circumstances can lead to even more radical effects, such as a change in the set of possible outcomes to be considered. A coin tossed on a rough table, for example, might land upon its edge with some positive probability, an impossible event if the coin is tossed on a smooth table. The outcome space of the former must therefore be specified in a different manner than the outcome space of the latter. See Gillies (2000, p. 812).

[14] Another way of putting this is that lotteries have their probability assignments *ceteris paribus*, as further information may impact them or even disqualify them as lotteries. But this is no different from the way that

these features are unknown. This is a good thing, too; if *every* potentially relevant feature of a coin toss were known, then presumably the outcome of the toss could be predicted with certainty, and it would not count as a lottery at all. Lotteries thus live in a zone in which it is always possible that further information relevant to prediction may be forthcoming; and if some of that information arrives, a lottery may have to be redescribed or even excluded altogether from the class of lotteries.[15]

It is possible to make one further distinction between types of lotteries, this time within the class of lotteries with risk. One obvious distinction to draw is between lotteries which generate all outcomes with *equal* probability from those that do not. The latter are "weighted" to favor some outcomes more than others, although not necessarily by design. For this reason, I shall distinguish within the category of lotteries with risk between *fair lotteries*, which are lotteries with risk which generate all outcomes with equal probability, and *weighted lotteries*. The former term captures nicely the attractiveness of such lotteries in a wide variety of circumstances. Indeed, intuitively most of the uses to which one might put a lottery call for a fair lottery.[16] But my intention here is merely to distinguish between different types of lotteries, not to argue for using particular types in particular contexts. A full-scale treatment of the latter topic is beyond the scope of this paper, although I shall say a few things about it in the final chapter.

---

a process with certainty must be described as certain only *ceteris paribus*, as there may be some as-yet unknown factor that must be present for the outcome to be generated with certainty.

[15] This need not mean that such information is equally likely to be forthcoming for all lotteries. It might be much easier, for example, to discover new information about a lottery with uncertainty, about which almost nothing is known, than a lottery with risk, about which enough is known to make probability assignments. But further revelations are always theoretically possible regarding any lottery.

[16] In their discussion of lotteries and justice, Lewis Kornhauser and Lawrence Sager take for granted that equiprobability is the property that unifies all paradigmatic cases of lotteries (coin tosses, etc.). See Kornhauser and Sager (1988, p. 485). They thus equate the class of lotteries with the class of fair lotteries. Given that it is clear that, for the purposes they wish to achieve, only fair lotteries are relevant, this is an understandable equation. It is also made by many other scholars concerned with lotteries. Such an equation, however, makes it difficult to classify processes that resemble the paradigmatic cases except with regard to their lack of equiprobability.

5. Lotteries and Decision-Making: The purpose of a lottery is to make decisions. Its appropriateness for this purpose is an open question; clearly, not *all* decisions should be made by lot. But assuming that some decisions are legitimately made in this manner, it is worth briefly laying out the basic steps necessary to decision-making processes and the role that lotteries are capable of playing in such processes.

All decision-making processes consist of two steps. Jon Elster describes these two steps as follows:

> A general theory of human action…can be sketched as follows. To explain why a person in a given situation behaves in one way rather than another, we can see his action as the result of two successive filtering processes. The first has the effect of limiting the set of abstractly possible actions to the *feasible* set, i.e. the set of actions that satisfy simultaneously a number of physical, technical, economic and politico-legal constraints. The second has the effect of singling out one member of the feasible set as the action which is to be carried out (Elster's emphasis; Elster, 1984, p. 76).

Whenever an actor decides how to act, that actor's behavior must be determined by at least one filter of each sort (cf. Føllesdal, 1982, pp. 306-307). Both filtering operations may be performed well or poorly. A decision-maker might overlook options worthy of consideration, or waste time considering options that are unfeasible or otherwise worthless. And obviously, she might select the *wrong* feasible option at the end of the day. (There is no way to say what "wrong" means here without specifying the type of decision to be made.) But right or wrong, both filtering options must be performed.

Note that an actor might have to engage in this two-step process more than once. It might happen, for example, that the actor applies her second filter to the feasible set of options and finds that more than one option remains. This might happen, for example, if the filtering process is supposed to eliminate all options except the one that maximizes some criterion (e.g., happiness or utility), and more than one option accomplishes this end (i.e., there is a tie). In such a case, the agent must repeat the same two steps. The outcomes yielded by the second stage the first time become the input for the first stage the second time. Of course, the filtering processes used the second time around will not be identical to those used the first time around; if they were, then further filtering would be impossible. Nevertheless, the two filters used the second time around play the same sort of role as those used in the first instance.

Obviously, there are a near infinite number of ways that an agent could either generate a set of feasible options or select a single option from such a set. The lottery constitutes one means of accomplishing the latter task. How does it perform such a task? The process works as follows:

1) The first filter (of whatever kind) generates a list of feasible options for consideration.
2) The actor selects a lottery capable of generating the same number of outcomes as are contained in the set of feasible options to consider.
3) Each outcome of the lottery is matched to a different filtering process that weeds out every option except one. In other words, every outcome of the lottery *x* is matched to a filter of the form, “Filter out all options except *y*,” such that every outcome *x* and every option *y* is used once and only once.

4) The lottery takes place, and an outcome is generated;
5) The actor employs the filter that was matched to the outcome generated by the lottery. The actor then takes the action that remains after employing this filter.

The lottery thus functions as a "second-order" filter of the second type; it selects a second filter for the actor to employ.[17] This selection takes place by matching each possible filter to a different outcome of the lottery, and then generating an outcome.

Suppose, for example, that some agent must decide which of two claimants, Annie and Howard, is to receive some good. Suppose further that this agent makes this decision via a coin toss. The agent accomplishes this task by (say) matching the lottery outcome "heads" with the filter "select Annie over Howard" and the outcome "tails" with the filter "select Howard over Annie." Once the coin is tossed, the agent takes the action dictated by the filter that was matched with the outcome generated—heads, give the good to Annie, tails, give the good to Howard. (One might say loosely that the outcome "heads" was matched with Annie, and the outcome "tails" was matched with Howard.)

This same logic is at work when other kinds of lotteries are employed, although this fact may not be obvious at first glance. Suppose, for example, that instead of tossing a coin, the agent had decided to ask Annie and Howard to draw straws, with the long straw winning the good. If one takes the set of possible outcomes of the lottery to be "long straw" and "short straw," then it is unclear how one might match these outcomes to possible filters before the lottery is employed. But in fact the set of outcomes yielded by this lottery consists of "long straw to Annie, short straw to Howard" and "short straw to

[17] It is thus wrong to suggest, as Isaac Levi does, that employing a lottery to make a decision in effect introduces a new option to the set of feasible options. See Levi (1986, p. 72).

Annie, long straw to Howard."[18] Each of these outcomes could clearly be matched to a filter in the same manner as the coin toss described above.

6. The Constructive Use of Lotteries: In order for a lottery to play the decision-making role described above, it must possess all three of the properties specified in its definition. Obviously, it must be a process capable of giving rise to outcomes. Obviously, the process must give rise to a discrete number of outcomes—specifically, a number of outcomes equal to the number of options from which a decision is to be made. But most importantly, the outcome of the lottery must not be predictable with certainty in advance. This last point is important enough to merit further exploration here.

Once the outcome of the lottery is known, its use to make a decision is a simple exercise in modus ponens. If outcome *x* arises, option *y* is to be selected (i.e., all options except *y* are to be filtered out); outcome *x* has arisen; therefore, option *y* is to be selected. In this respect, lotteries resemble other processes that could be used to make decisions—e.g., a decision-making rule such as utilitarianism. The utilitarian, after all, makes similar use of modus ponens. If option *x* yields greater utility than option *y*, select option *x* over option *y* (i.e., filter out option *y* if option *x* is available); option *x* does yield greater utility than option *y*; therefore, select option *x* over option *y*. But there is a critical distinction between lotteries and other processes. An agent employing a utilitarian decision-making rule can employ that rule regardless of whether or not the option yielding the greatest utility is known in advance. But an agent who seeks to employ a lottery to make the

[18] One could abbreviate this simply as "long straw to Annie" and "long straw to Howard." Once the claimant drawing the long straw is identified, the outcome of the entire lottery is uniquely identified given that only one long straw is to be drawn (and, by assumption, there are no relevant distinctions between the short straws). In the two-option case, knowing which claimant received the short straw uniquely identifies which claimant received the long straw, and thus which filter is to be employed. But this will not be the case when there are three or more claimants drawing straws. In such a situation, knowing that a claimant has drawn the short straw eliminates one possible outcome of the lottery, but does not uniquely identify the outcome.

decision cannot know the outcome in advance; if she does, she is simply not using a lottery. If the outcome of a lottery were somehow known in advance when a decision was made with it, an essential property of that lottery—one of the properties that distinguish it from non-lotteries—would not have been used in making that decision. The lottery might have been used, but it would not have been used *as a lottery*.

This point is important, as there are many processes that resemble lotteries in some respects but are not *employed as lotteries* in decision-making processes. Consider an election for a particular office, for example. (I owe this example to Eric MacGilvray.) An election is a process that yields a variety of outcomes, each of which consists of a different tally of votes. The outcome of the election is typically not known with certainty.[19] And each outcome of the election is matched to a different option—specifically, a different candidate taking office. To this extent, an election resembles a lottery. But suppose each voter had truthfully announced in advance how he was planning to vote. In such a case, the outcome of the election would be known in advance. But this would not in any way invalidate the appropriateness of relying on an election to fill the office. This means that the lack of predictability is not essential to the use of election. The reasons for using an election make no essential reference to the lack of predictability of the outcome; instead, they may refer to the need to identify the best candidate, or to ensure that the political system enjoys legitimacy, or to treat all voters as equals, or the like.[20] None of these factors require unpredictability, and therefore none point to the use of a lottery.

---

[19] Cf. Eugene Meehan (1988, p. 133): "Actions of such collective bodies [as bodies of voters] come very close to being in the same class with hurricanes and floods, 'acts of nature' of a peculiar kind."

[20] Brian Barry offers an argument for democratic elections that constitutes a partial exception to this rule. Barry claims that "The most important point about a system of election for representatives is that it

The same holds true of procedures that allocate goods on the basis of determinate criteria such as need or merit or willingness to pay (as in an auction). In many such cases, it might be foreseen well in advance who is the neediest, or the most meritorious, or the richest. But this does not obviate the use of such a procedure to make the decision. And this is a clear sign that, even in cases in which the neediest candidate, say, is not known in advance, the process is not being used as a lottery. There is a clear test as to whether this is the case. If the agent making the decision were unable to substitute a suitably-described paradigmatic lottery—a coin toss, a die roll, or the drawing of straws—for the procedure of selection by merit (or whatever), then regardless of the predictability of that procedure, it is not being used as a lottery.[21]

It is possible to use even a paradigmatic case of a lottery in a manner that does not treat it as a lottery. Suppose, for example, that a deeply religious individual believes that the casting of lots reveals the will of God. (This belief was quite widespread in Western antiquity—there are numerous uses of the lot described in the Bible—and it persisted into the modern era via the Moravian Church.) Such an individual might use a lottery to decide what to do in a given context, but she does not use the lottery as a lottery. Instead, she uses it purely as an information-revelation procedure. Had she some other means of ascertaining the will of God—a telephone to heaven, for example—she would use that instead. And if God had spelled out in advance what He wanted done, there would be no

---

provides an intelligible and determinate answer to the question why these particular people, rather than others perhaps equally well or better qualified, should run the country" (Barry, 1979, p. 193). Were this argument correct, it would be unclear whether an election functions in any way other than as a lottery. After all, a coin toss is perfectly capable of providing an answer to that same question, an answer that (barring further argument) seems just as intelligible and determinate as counting noses.

[21] Kornhauser and Sager (1988, p. 511, n. 1) note that when an adjudicatory mechanism is used, "the distribution of the prize among the candidates is uncertain. In this sense every adjudicatory mechanism looks, in part, like a lottery." But they do not spell out what, given this resemblance, distinguishes a lottery from an adjudicatory mechanism that relies upon need, merit, or whatever.

reason to appeal to any such procedure. The lack of predictability plays no role at all. (The same holds true of an individual who believes that "luck" or "fate" or some other superhuman force controls the outcome of the lottery.)

In many cases, it might not be apparent whether a particular process is being used as a lottery or not. Consider, for example, a rather unusual lottery described by George Sher. Sher imagines an agent, *m*, who must decide which of claimants *n*, *o*, and *p* are to receive a particular good *G*. He envisions the following procedure:

> We can imagine a case in which *n*, *o*, and *p* are persons of unknown racial and religious backgrounds, and in which *m* decrees that *G* will be awarded to whichever one of them is discovered, through a genealogical search, to have the fewest Jewish ancestors (Sher, 1980, p. 206).

There are at least three distinct ways in which this "fewest Jewish ancestors" rule might be employed in a decision-making process.

1) If *m* is anti-Semitic, he might regard Jewish ancestry as a reason against giving *G* to someone. In this case, Jewish ancestry might be one among several factors (such as need) that *m* considers in deciding who is to receive *G* (i.e., who is not to be filtered out). In that case, Jewish ancestry is not being employed by *m* as a lottery; if *m* knew in advance how many Jewish ancestors *n*, *o*, and *p* each had, it would make no difference in his employment of this criterion.

2) It could be that *m*, though anti-Semitic, recognizes that some other criterion, such as need, is the legitimate means for distributing *G*, regardless of racial or ethnic background. But *m* might believe that, if *n*, *o*, and *p* are equally needy, such that a tie is generated between them, Jewish ancestry (or lack thereof) may

legitimately be employed as a tiebreaker. Here Jewish ancestry in effect becomes a second second filter, to be used after the first second filter (need) failed to yield a unique decision. Here *m* would be using a lexicographic procedure to decide who receives *G*, just as an agent employing Rawls' conception of justice invokes the Difference Principle to decide between economic systems only after other principles (Equal Basic Liberty and Fair Equality of Opportunity) have proven equally satisfied by those systems. Here again Jewish ancestry is not functioning as a lottery, for the same reasons as in case #1.

3) Finally, *m* might genuinely envision using Jewish ancestry to constitute a lottery. In such a case, he should be indifferent between using the "fewest Jewish ancestors" lottery and a different lottery, such as asking *n*, *o*, and *p* to draw straws.

If *m* proposed the use of a "fewest Jewish ancestors" lottery, then *n*, *o*, and *p* might be legitimately unsure as to *m*'s motivations. This uncertainty might in itself constitute a reason for using a different lottery, assuming that all agree that a lottery of some sort is required.

In short, when a process' outcome is unpredictable, but its unpredictability makes no difference to its employment in making a decision, that process is capable of being *described* as a lottery, but it is not being *used* as a lottery in making the decision in question. The reasons for using the process are not reasons for using a lottery.

None of this answers the question, of course, of what are the reasons for using a lottery. The final section of this paper provides some tentative reflections on this question.

7. Conclusion: When a lottery is used in decision-making *as a lottery*, all of its essential properties are used, including the unpredictability of its outcome. Therefore, if a lottery is being used *qua* lottery to make a decision, changing its lack of predictability should change its desirability in that decision-making scenario. Moreover, there should be other lotteries—most likely the paradigmatic cases, such as a coin toss—that can be substituted without difficulty in making that decision. If difficulties arise in making such a substitution, that is a clear sign that the reasons for using the process are not reasons for using a lottery.

All of this implies that, if there are reasons for using a lottery, those reasons must depend upon all of the properties of a lottery, especially the unpredictable nature of its outcome. What sort of reasons might work this way? I suggest that any argument for using a lottery for decision-making must rest upon the *sanitizing effect* that an unpredictable process can have upon decision-making. Decisions, after all, can be made for good or for bad reasons. A lottery, however, selects an outcome based upon *no* reasons, good or bad. If there is an argument for using a lottery, therefore, it must depend upon the ability of the lottery to avoid making use of bad reasons. The decision, then, must be of the sort that there is the danger that bad reasons may come into play, and yet there is minimal danger of overlooking good reasons.[22]

There are, I believe, many decisions that might satisfy this condition. Consider, for example, the problem of allocating goods in accordance with some conception of justice. That conception might require that need, merit, desert, expected utility, or other

[22] Of course, if an agent makes a decision via lottery, there will be a reason of sorts for that agent's final decision. "Why did you select option *x*?" "Because the coin toss came up heads." But this is a second-order reason, not the sort of reason one might expect to play a role in a decision. A lottery is thus most useful when one has second-order reasons to avoid using first-order reasons, because of the fear that bad first-order reasons may determine the decision.

factors be used to decide which of several individuals is to receive some good (or exemption from some bad, such as military conscription). The conception thus provides reasons for favoring one candidate over another. But it may often be the case that the reasons favoring two or more individuals are equally good. In such a case, there would be a tie among some of the candidates that must be resolved in order for the allocation to take place. (This is a case in which the second filter employed in decision-making failed to yield a unique outcome.) One might argue that justice demands that no candidate be awarded the good on the basis of any reasons except those specified by the conception of justice. But if those reasons are exhausted, and yet there is a need for further discrimination, the only way to make the decision that is compatible with the demands of justice is to break the tie on the basis of no reasons at all. And only a lottery can accomplish this task.[23] Other arguments for the use of lotteries, I suggest, must have essentially the same form.

Knowing the form that arguments for the lot must take makes it possible to revisit the three different types of lotteries described above. The sanitizing effect of lotteries is maximally achieved when a fair lottery—a lottery in which all outcomes happen with equiprobability—is employed (cf. Kornhauser and Sager, 1988, p. 488). A lottery that favors some outcomes over others—what I have called a weighted lottery—will, when incorporated into decision-making, ensure that some options are selected more often than others. This raises an obvious question—which options should the lottery favor? There is no way to answer this question without listing reasons why some options should be chosen over others, but creating such a list and making it the basis for a decision is precisely what the lottery is used to prevent. This suggests that an argument for the use of

[23] For a more elaborate version of this argument, see Stone (forthcoming).

weighted lotteries will have to take the form of a tradeoff between two considerations, one pointing to the use of a fair lottery, the other pointing away from it.[24]

It also suggests that there is little, if any, reason ever to use a lottery with uncertainty. Why use a lottery whose nature is fundamentally unknown, when lotteries exist with well-understood properties? It is hard to imagine an argument for a lottery that did not indicate a desire either for equally likely outcomes, or for some outcomes to be chosen more often than others. A lottery with uncertainty can guarantee neither. While any case for making decisions by lot must, I have argued, make essential reference to our ignorance of the lottery in some respect, not all forms of ignorance lead to bliss.

One final point about lotteries is worth making here. I have presumed throughout that the agent using a lottery in decision-making could be either an individual or a collective entity. Much of the argument offered here, however, becomes much more complicated when a collective agent (such as a nation or other political unit) is involved. For example, my definition of a lottery focuses on the inability of the decision-making agent, given the evidence available to her, to predict the outcome of the lottery. This is fairly straightforward when the agent is an individual, but what does it mean to speak of a collective body's ability to predict the future? What does it mean to speak of the evidence available to that body, or the correct amount of evidence that body should gather? Any satisfactory theory of the role of lotteries in collective decision-making must confront questions such as these.

[24] An example of such a tradeoff can be seen in the decision made by the National Basketball Association (NBA) to allocate draft picks in an order determined via a weighted lottery, with more weight being given to the worst-performing teams. The NBA had to balance the desire to strengthen weaker teams against the fear of creating bad incentive effects (i.e., a team might purposely lose games in order to increase its chance of getting excellent draft picks.)

Bibliography

Ayer, A.J. “Two Notes on Probability.” In *The Concept of a Person and Other Essays*. London: Macmillan, 1963.

Barry, Brian. “Is Democracy Special?” In Peter Laslett and James Fishkin, eds. *Philosophy, Politics and Society*. Fifth Series. New Haven: YUP, 1979.

Duxbury, Neil. *Random Justice: On Lotteries and Legal Decision-Making*. Oxford: Clarendon Press, 1999.

Elster, Jon. *Ulysses and the Sirens*. rev. ed. New York: Cambridge University Press, 1984.

Elster, Jon. *Solomonic Judgments*. New York: Cambridge University Press, 1989.

Føllesdal, Dagfinn. “The Status of Rationality Assumptions in Interpretation and in the Explanation of Action.” *Dialectica* 36:4 (1982): 301-316.

Gillies, Donald. “Varieties of Propensity.” *British Journal of the Philosophy of Science* 51 (2000): pp. 807-835.

Goodwin, Barbara. *Justice by Lottery*. 2nd ed. Charlottesville, VA: Imprint Academic, 2005.

Kornhauser, Lewis A. and Sager, Lawrence G. “Just Lotteries.” *Rationality and Society* 27 (1988): 483-516.

Knight, Frank H. *Risk, Uncertainty and Profit*. Washington: Beard Books, 2002.

Levi, Isaac. *Hard Choices: Decision Making under Unresolved Conflict*. New York: Cambridge University Press, 1986.

Meehan, Eugene J. *The Thinking Game: A Guide to Effective Study*. Chatham, NJ: Chatham House Publishers, 1988.

Rescher, Nicholas. “Choice without Preference.” *Kant-Studien* 51 (1959-1960): 142-175.

Sher, George. “What Makes a Lottery Fair?” *Noûs* 14 (1980): 203-216.

Stone, Peter. “Why Lotteries Are Just.” *Journal of Political Philosophy*, forthcoming.

Waldron, Jeremy. “Legislation and Moral Neutrality.” In Robert E. Goodin and Andrew Reeve, eds. *Liberal Neutrality*. New York: Routledge, 1989.