

Doing the best one can: a new justification for the use of lotteries

ITTAY NISSAN-ROZEN

Hebrew University in Jerusalem

Abstract: In some cases in which rational and moral agents experience moral uncertainty, they are unable to assign exact degrees of moral value—in a non-arbitrary way—to some of the different acts available to them, and so are unable to choose with certainty the best act. This article presents a new justification for the use of lotteries in this kind of situation. It is argued that sometimes the only rational thing for a morally motivated agent to do here is to use a lottery.

Keywords: moral uncertainty, lotteries, indivisible goods

JEL Classification: B41, D01, D03, D80

Many people share the intuition that in some choice situations using a lottery among (some of) the acts available to an agent is the morally right thing to do. In the philosophical literature several justifications for this intuition are presented. The most famous is John Broome's justification, which is based on the idea that what makes using a lottery the morally right thing to do (when it is the morally right thing to do) is that it is fairer than any of the definite choices available to the agent.¹ Thus, Broome's explanation of what makes a lottery right has two parts: first he presents an account of the fairness of lotteries, and second he argues that in some situations the fairness consideration is strong enough to make the fair act the right act.

In this paper I will present a new justification for the rightness (in a sense specified below) of lotteries. According to my justification a lottery is justified in some situations where an agent suffers from moral uncertainty, i.e., in some situations when an agent is unsure what the morally right thing to do is (however, not in every situation in which this

¹ See, e.g., Broome 1990, 1991, 1994. Other discussions of the questions include Hooker 2005; Sher 1980; Saunders 2009; Rescher 1969; Glover 1977.

is the case). I will argue that in these situations using a lottery is the best one can do, given one's moral uncertainty. I will also characterize the set of situations in which a lottery is justified according to my account and present an explication for the term "the best one can do".

However, unlike Broome, I will not argue that using a lottery, when it is the right thing to do according to my account, is also the fair thing to do. One could take a further step and try to argue that what makes a lottery right according to my account is also what makes it fair. Hence, one could argue that being fair is just doing the best one can to do the right thing. I think there might be good reasons to take this further step,² but I will not argue for it here. Here I only present a justification for the use of lotteries, not an account of fairness.

Is my account a rival to Broome's account? Not necessarily and actually there are reasons to think that the two are never rival accounts, as they seem to address different issues. One can at the same time hold the position that some lotteries are morally right for the reasons Broome presents and that some are right (in a different sense of "rightness") for the reasons I present (and that some may be right for other reasons).

The rest of the paper will be organized in the following way. In the first section, using Broome's discussion of the tension between the sure-thing principle and the rightness of lotteries, I will present some background issues that will be of later use (I will not, however, present Broome's account). In the second section, I will discuss the idea of moral uncertainty; and in the third section, I will present my account for the rightness of lotteries using a formal framework. In the fourth section, I will discuss the recommendations that my account gives in some cases and will argue that these provide independent support for my account.

THE FAIRNESS OF LOTTERIES

Broome's starting point is the intuition that: "Sometimes a lottery is the fairest way of distributing a good" (Broome 1990, 87). Broome also holds that because of this fact "there will certainly be some

² Hooker (2005) acknowledges (and refers to others who acknowledge) that "fair is often used with a very broad meaning. A 'fair decision', in this very broad sense of 'fair', means a decision that appropriately accommodates all applicable moral distinctions and reasons" (Hooker 2005, 331). This understanding of the term "fair" is in line with the "being fair as doing the best one can" thesis, only that under the explication presented here for "doing the best one can" such an understanding of fairness can also explain why lotteries are sometimes fair.

circumstances where it is better to hold a lottery than to choose the best candidate deliberately” (Broome 1990, 99).

This latter claim, poses a problem for Broome that he has to deal with even before presenting his justification for the intuition he started with: it seems that any moral preferences ordering that ranks a lottery between two actions above both of these actions must violate an intuitive principle of rationality called the sure-thing principle (SP). The sure-thing principle requires that when an agent is uncertain what the consequences of some of the actions available to him will be, then when he evaluates these actions he can disregard any state of the world in which all of them bring about the same outcome, as shown in Table 1:

Table 1

	ω_1	ω_2
L	A	B
A	A	A
B	B	B

The SP requires that if the agent prefers act A to act B then he should prefer act A to act L and act L to act B. Thus, it is easy to see that a lottery between two alternatives should never be preferred to both of them.

One way to deal with this problem is to reject the SP in moral contexts.³ However, this is not the strategy Broome adopts and he (as well as others) has presented very convincing arguments against it (see Broome 1984, section 2). To deal with this problem Broome suggests that in cases in which a lottery seems to be morally preferable to any of the alternatives over which it is defined, we have to include the fairness achieved by using the lottery in the description of the outcomes.⁴ By following this suggestion, the SP is not violated because it does not apply. This is illustrated in Table 2:

³ This is the position adopted, for example, by Diamond (1967), who first introduced this problem.

⁴ See Karni 1996, for a similar suggestion.

Table 2

	ω_1	ω_2
L	A achieved by a lottery.	B achieved by a lottery.
A	A achieved by a definite choice.	A achieved by a definite choice.
B	B achieved by a definite choice.	B achieved by a definite choice.

Since, under the new interpretation of the situation, the two possible outcomes that act L might bring are different from the outcomes acts A and B bring, the SP does not apply to the decision problem and so is not violated.⁵

However, regardless of the question of whether the SP must be violated in cases where a lottery is ranked above all of the acts over which it is defined, the following claim does hold: whenever an agent prefers a lottery to all the definite acts over which it is defined, the agent does not maximize the expectation of a quantity that we can call goodness.⁶ He may be maximizing expected moral value—i.e., the expectation of the overall value of an act, given both the act’s expected goodness and its fairness—but not expected goodness.

The last observation is the reason why I have lingered on the discussion above. The important distinction I want to make is between accounts—like Broome’s—that recommend a lottery even in cases in which it is possible for the agent to choose a definite act with higher expected goodness, and accounts that recommend a lottery only when it is impossible for the agent to do that. Accounts of the latter type can, for example, recommend a lottery in cases in which the agent is morally indifferent between two acts. In such cases, any lottery between the two

⁵ In the literature there are several objections to Broome’s use of the “redescribing the outcomes” strategy (see Steele 2006, for a good overview). Broome (1991) discusses one of them: the violation of the “rectangular field assumption”. I will not discuss these objections here. See Bradley 2007, however, for a formal framework that resolves Broome’s worries.

⁶ “Goodness”, as I use the term here, refers to the moral value of an outcome (i.e., an act in a specific state), *not including* the moral value added to the act that brings this outcome in virtue of it being a fair lottery. “Expected goodness” is to be understood in an analogous way. I think that Broome uses the term in the same way most of the time (at least in his earlier work), but as there are some places in which he makes comments that can be understood as implying the contrary, I do not want to argue that he does. In any case, this is the way I am going to use the term here.

acts has the same expected goodness as each of the two acts and so it is impossible for the agent to choose an act with higher expected goodness. However, such accounts cannot recommend a lottery in cases in which the agent morally prefers one definite act to another.

The account that I will put forward here can do—in an important sense—both: it never recommends a lottery when it is possible for the agent to choose a definite act with higher expected goodness, and it does sometimes recommend a lottery over all the definite acts over which it is defined, even when the agent is not morally indifferent between all of them. How can this be the case? The answer lies in my use of the idea of moral uncertainty. I will discuss the idea in more detail in the next section, but even before doing that, it is easy to see how using this idea can make such an account possible.

When an agent is certain that one of the acts available to him has higher expected goodness than the other(s), but is not certain which one it is, then although it is metaphysically possible for him to choose the act with the highest expected goodness (it is, after all, one of the acts available to him), it is not epistemically possible for him to do so. In some such cases (but not in all of them), my account will recommend a lottery. Thus, in those cases, a lottery will be recommended even though the agent is certain that there is another act available to him with higher expected goodness, and yet the requirement to always choose the act with the highest expected goodness (when it is epistemically possible to do that) is kept.

So my account recommends choosing a lottery—when it does recommend that—not because this is the “objectively” right thing to do, but rather because, given the agent’s epistemic state, this is the only “subjectively” right thing for him to do.

To avoid confusions, a terminological comment is necessary here. As mentioned—and as will be discussed at more length in the last section—my account is consistent both with the position according to which being a fair lottery adds some positive moral value to an act, and with the position according to which this is never the case. In order to avoid taking a stand in this debate, I need to use three different terms that describe different types of moral features of acts. I have already introduced the first one—“goodness”—explicitly: the “goodness” of an act is the moral value of the outcome that an act will actually deliver (not including the moral value added to it in virtue of it being a fair lottery).

Now we see that there are two more moral features of acts that will play a role in my discussion. The first one is the overall moral value of an act (including the possible moral value added to the act in virtue of it being a fair lottery). I will use the term “overall moral value” to refer to this feature. The second one is the possible moral value added to an act in virtue of it being a lottery. I will use the term “fairness” to refer to this feature.

When an agent does not suffer from moral uncertainty he ought to choose the act with the highest overall moral value. I have just used the term “the objectively right act” in order to refer to this act and will continue to do that. However, when an agent does suffer from moral uncertainty he cannot choose the objectively right act with certainty. Which act should the agent choose in such a case? Whatever the answer to this question is, we need a term in order to refer to this act. I have just used the term “the subjectively right act” in order to refer to it and I will continue to do that.

Notice that in light of the distinction between the best act and the right act, we can already distinguish between two types of moral uncertainty: regarding which act is the objectively right act to choose and regarding which act is the best act (when the “best act” is the one with the highest expected goodness). My discussion of moral uncertainty, in the next two sections, will only refer to the latter type of moral uncertainty, i.e., to uncertainty regarding which act is the morally best act (and for convenience I will just use the term “best act”).

It should be noted, though, that by limiting my discussion only to this type of moral uncertainty, I do not mean to suggest that Broome’s account—and others like it—are either false or cannot be applied to the cases that I discuss. I offer an alternative account for the use of lotteries that can justify choosing a lottery even when being a fair lottery does not add any moral value to an act (or adds some moral value but not enough to make the fair act, the right act), but this account can live peacefully with other accounts. In the last section, I will go back to the above distinction and will discuss the advantages that my account has in this context. I move now to discuss the idea of moral uncertainty.

MORAL UNCERTAINTY

The idea of being uncertain regarding what is the morally best thing to do is the key element that will help me develop my account. It is possible to make a distinction between three different types of

uncertainty regarding which act is the morally best act to choose in a given situation, only one of which will be dealt with here. The first type is that of such uncertainty that results exclusively from uncertainty regarding the truth of some non-normative claims. For example, an agent might be uncertain whether one act, A, is better than another act, B, because he is uncertain what the consequences of the acts would be. If the agent knew for sure what their consequences would be, he would not be uncertain regarding which act is better. There are some discussions in the literature (see, e.g., Diamond 1967) regarding the question of how a rational moral agent should make his decisions in the face of such uncertainty, but these will not concern us here.

The second type of uncertainty is that which cannot be reduced to uncertainty regarding the truth of some non-moral claim, but still can be reduced to uncertainty regarding the truth of some general moral claim. For example, an agent might be uncertain whether one act, A, is better than another act, B, only because he is uncertain whether consequentialism is true or not. If he was sure that consequentialism is true (or not) he would not be uncertain regarding which act is the best act to choose. This type of uncertainty is the one which is usually discussed in the literature under the title “moral uncertainty” (see, e.g., Lockhart 2000; and Sepielli 2009).

Most of the accounts that belong to this literature treat moral uncertainty much in the same way that decision theory treats uncertainty regarding states of the world, i.e., by requiring that in face of moral uncertainty one should maximise expected goodness (relative to the moral uncertainty one suffers from). Such a requirement, however, is based on two assumptions: 1) one is able to tell how good or how bad every possible act is, according to each of the moral theories (or general moral claims) one believes might be true; and 2) one is able to compare these values across theories.

Much of the discussion in the literature is dedicated to an examination of the second assumption, i.e., to the question of which principles should govern the inter-theoretical comparison of moral value. However, in some cases, prior to dealing with the second assumption, one has to deal with the first one. Before one can present a prescription for how to compare the degrees of moral value different moral theories assign to an act, one has to explain how it is possible to get those degrees according to a single theory.

While in some choice situations it seems that the theories themselves might offer an answer to this question, in other choice situations this cannot be the case. Situations of the latter kind are usually (but maybe not always) those in which different moral considerations (according to the same theory) push in different directions and thus, in order to assign exact levels of goodness to the acts, the theory must assign weights to the different considerations.

We can judge, for example, that saving the life of another person is better than slightly improving his wellbeing, but it is really hard for us to say exactly how much better it is. Thus, when we have to decide between saving the life of one person and slightly improving the wellbeing of many people we may become uncertain regarding what we ought to do in a specific case (i.e., when the number of people whose wellbeing we can improve is high enough). The same is true for the moral judgements we make, under the assumption that a specific moral theory (or a general moral claim) is the right one. One can accept, for example, utilitarianism, or even a specific type of utilitarianism, but still be uncertain regarding how different aspects that increase utility ought to be weighed against each other.

Notice that in this kind of case the moral uncertainty arises *because* we do not have direct access to degrees of moral value. If we had such access we would know how to weigh the different moral dimensions against each other and would just form our moral preferences according to the levels of expected goodness of the different acts.

Now, one might argue that although an agent does not have direct access to the level of goodness of the different acts available to him; he does have direct access to the level of goodness of the different acts conditional on some betterness relation holding between them. In other words, it might be that an agent who is uncertain whether act B is better than act C, but is certain that act A is better than both act B and act C, is also certain, *regarding every possible lottery between A and C*, that if it is the case that act B is better than act C, this lottery is either better or worse than act B. This is what it means, for a rational agent, to have direct access to the level of goodness of the acts conditional on some betterness relation holding between them.

I do not want to argue that such cases never happen. I certainly believe that in many cases, people have partial information regarding degrees of goodness conditional on some betterness relation holding.

However, I also believe that there are cases in which this type of information is unavailable.

Consider an agent who tries to follow the strategy described above when coming to make a moral decision between two acts, of which she is uncertain which is better than which. The agent formulates a number of hypotheses regarding the degrees of moral value of the acts available to her, conditional on either of the two betterness relations holding. Then she has to assign a probability value to each one of these hypotheses. On what basis can this be done? Surely we do not want to argue that she should do this arbitrarily. Arguing this is like arguing that she should choose an act arbitrarily as by assigning different probabilities to the hypotheses, she can make either of the acts the one that maximises expected goodness.

A more plausible answer is that she should do this according to her actual degrees of beliefs in these hypotheses, which should probably be based on what she takes to be moral evidence. However, it is not clear what can constitute evidence for a specific hypothesis regarding the exact degree of moral value of an act, other than the kind of betterness judgements that assigning such a degree to the act (together with assigning other degrees to other acts) leads to. Actual moral theories do not give us exact degrees of moral values in the kind of situations I am referring to and people do not generally have intuitions regarding such levels. What people do when they have to assign such levels is to implicitly judge which one is better among many possible acts.

Thus, in such choice situations, since moral agents do not have direct access to the degrees of moral value that different theories assign to different acts, the need arises to present an account for decision making under conditions of moral uncertainty that does not make any reference to these degrees of moral value. This is not because the requirement to maximize expected moral value is not a principle of rationality. I believe it is. Rather, it is because when one does not have access to these degrees, one cannot possibly obey this requirement and thus the need arises for an alternative decision rule.

This last type of moral uncertainty, i.e., moral uncertainty that cannot be reduced either to uncertainty regarding non-moral claims or to uncertainty regarding general moral claims, is the one which I am going to discuss in the next section.

MORAL UNCERTAINTY AND LOTTERIES

When an agent is uncertain regarding the morally best thing to do in a specific situation because she does not have access to the degrees of moral value of the different acts available to her, but still must make a decision, what should she do? The immediate answer is, I think, that she should try to minimize this uncertainty as much as she can: she should spend some time reflecting on the matter, she should consult with people whose opinions she values, she should read some books, and so on. But when she is done with this process, when she has used any sources of moral information available to her, then if she is still uncertain regarding what is the best thing to do, she has no plausible alternative but to go with the judgements she is more certain about. So if she believes that some act A is better than another act B more strongly than she believes that B is better than A, she should choose A over B. Let us call this requirement the Likelihood of Betterness Constraint (LBC).

Such a prescription clearly falls short of the ideal of maximizing expected moral value (which the agent cannot obey because she does not have access to degrees of moral value), but it does require that the agent make use of the information she does have, i.e., her degrees of belief regarding the moral betterness relations that hold between different acts.⁷

Notice, however, that if one accepts the LBC, but still wants one's moral choices to be transitive, one commits oneself to the requirement that for any three alternatives, A, B and C, if one believes it is more likely than not that A is better than B and that it is more likely than not that B is better than C, one must believe that it is more likely than

⁷ It might be argued that there are some situations in which the agent does not even have enough information to allow her to assign in a non-arbitrary way degrees of belief to the different possible betterness relations that might hold between different acts. Maybe there are such cases. However, there are certainly cases in which we do feel that we have enough moral evidence to assign in a non-arbitrary way degrees of belief to the different possible betterness relations, while we do not have enough moral evidence to assign in a non-arbitrary way degrees of beliefs to the different possible hypotheses regarding the exact level of moral value of each possible act. The reason for that is the one discussed in the previous section: in situations in which more than one morally significant aspect is involved, and thus the need to weigh the different aspects against each other arises, we usually do that using the betterness relations that we take to hold between different acts. See footnote 11, for a formal illustration of this point: the information that we need in order to assign degrees of belief to different betterness relations holding is strictly weaker than the information that we need in order to assign degrees of belief to different hypotheses regarding the degrees of moral value of different acts.

not that A is better than C. This condition does not follow, however, from the requirement that the agent believes with probability 1 that the betterness relation is transitive, as there are many probability distributions over the set of all possible orderings of alternatives that do not respect it.⁸

When this happens and one finds oneself in a situation in which in order to obey the LBC one must violate transitivity, at least one of these two requirements has to go. It is tempting to argue that this should be the LBC. This is so since it is clear that in such situations, by obeying the LBC, the agent will necessarily find himself violating a second-order moral judgement he should hold: the moral judgement that he ought to choose consistently when making moral choices.⁹ He will be violating this judgement by violating transitivity. Does this consideration give us a reason to reject the LBC? Not on its own, I will argue now.

Implicit in the objection to the LBC presented in the last paragraph is the assumption that the moral judgement that one ought to always choose consistently ought always to have priority over any other moral judgement. However, this assumption is dubious. The rationality requirements get their normative force from our belief that rationality is a guide for choices that will best serve the agent’s interests (in our case, moral interests: the interests of the agent when acting as a moral agent). The rationality here is instrumental rationality: there is no substantive moral value in obeying its requirements. The moral value of obeying its requirements comes from the further belief that doing so will best serve other purposes that do have intrinsic moral value.

⁸ Think of any “Condorcet paradox” style situation. For example, consider the following table:

1/3	1/3	1/3
A	C	B
B	A	C
C	B	A

Here, the agent believes with probability 2/3 that A is better than B, that B is better than C, and that C is better than A, and still gives a positive probability only to transitive rankings.

⁹ I do not argue that this is a moral judgement the agent should hold. I assume this for the sake of argument as I believe it is not an unreasonable position and is indeed a position to which many scholars are committed. John Harsanyi argued, for example, that “an individual making a moral value judgment must follow, if possible, even higher standards of rationality than an individual merely pursuing his personal interests” (Harsanyi 1978, 226). The argument that follows aims at showing that, even if one accepts the second-order moral judgement that one’s moral judgements ought to be consistent, in the kind of cases that I discuss here this moral judgment is defeated by other considerations.

However, when an agent believes it is more likely that one act is better than another rather than *vice versa*, it is clear that what will best serve the agent's moral interests—in the absence of sufficient information about degrees of moral value—is to choose this act over the other. Requiring that such an agent does otherwise, in the name of transitivity, amounts to putting the cart before the horse. It amounts to requiring that the agent gives priority to a moral judgement that gets its moral force from more fundamental moral judgements over one of those more fundamental moral judgements.

One might argue that, from a wider perspective, giving such a priority is justified, since by choosing in an intransitive way, the agent exposes himself to “money pumps”, or in the moral context, to “positive moral value pumps”. However, this argument misses the point. If the agent has good reasons, *in a particular case*, to suspect that by choosing intransitively, he will be drawn into a money pump, then this consideration ought already to be taken into account through his assessment of the possible consequences of the acts available to him. However, the mere possibility of being money pumped, *without having any reason to suspect that this possibility will actually be realised*, should not matter much to an agent who must make a specific decision.

So I think it is not the LBC that has to go in such situations. However, by relaxing transitivity we are left without a decision rule. Which act should the agent choose in a case where he has intransitive moral preferences over three acts and all three acts are available to him?

Here is one possible answer: if we allow the agent to use mixed strategies, i.e., if we require that the set of acts available to the agent is convex then—under a natural condition that describes the connection between the agent's factual beliefs and the agent's moral beliefs—there always exists an act that the agent believes is more likely or equally likely better than any other act available to him. In other words, there exists an act such that the agent believes that no other act is better than it. It seems reasonable to require from the agent to choose such an act.

This requirement can be seen as a generalisation of the requirement of not to choose an act to which another act is preferred, which is usually used to justify the transitivity axiom (e.g., Davidson, et al. 1955). When it is impossible for the agent to have transitive preferences, and when there is no act that is preferred to all the acts over which the intransitivity occurs, then this requirement cannot be respected. However, I will now show that the generalization of this requirement

(i.e., *the requirement not to choose an act if there is another act that one believes it is, more likely than not, better than the first*) can always be satisfied if the set of acts is convex, provided that one natural condition holds.¹⁰

Here is the condition. First informally: the agent's degree of belief that one mixed act is better than another equals his expected degree of belief that this act is better than the other. In order to express this condition formally, we need some more structure.

Let $\Omega = \{\omega_1 \dots \omega_n\}$ be a finite set of possible states. Let p be a probability distribution over Ω . Let $D = \{A, B, C \dots\}$ be a set of outcomes and let $E = \{a_1 \dots a_k\}$ be a set of acts, where an act is a function from Ω to D . Let \geq be the agent's moral preference relation over E . In addition let $>^*$ denote the moral betterness relation between pairs of acts, i.e., $>^*$ is a binary relation over elements of E . For simplicity, we will assume that for any two elements, a_i and a_j , $a_i >^* a_j$ or $a_j >^* a_i$. By assuming this, I am ignoring here the possibility that the agent gives a positive probability to the possibility that two acts are equally good, i.e., that neither one of them is better than the other. This assumption will make the discussion simpler and nothing is dependent on it.

Since we want to allow the agent to have beliefs regarding the betterness relation, we will usually need to refer to the betterness relation as a variable. In these cases we will just use the notation " $>$ ". Finally, let q be a probability distribution over all possible $>^*$ s. To be clear, the expression $q(a_i > a_j)$ denotes the sum of the probabilities q gives to all $>^*$ such that $a_i >^* a_j$.¹¹

¹⁰ There are certainly other decision rules one might consider as plausible candidates in this respect. See, for example, Loomes and Sugden's (1987) discussion of a similar question arising in the context of regret theory and SSB utility theory, and some of the articles in Arrow, Sen, and Suzumura's (2002) edited volume, which discuss a similar question in the context of social choice. I am not going to argue against any of these possibilities. Even if rationality does not require one to choose according to the decision rule I have offered above—which is in the spirit of Laffond, Laslier, and Le Breton's (1993) solution in the context of social choice—it would be hard to deny that rationality allows for using it. The discussion that follows in this section and the next one provides further support for this decision rule: by accepting it one gains a strong justification for the use of lotteries that does better—in terms of its ability to predict our moral intuitions regarding specific lotteries—than any other justification around.

¹¹ It is important to stress that by taking q to be a probability distribution over the set of all possible betterness relations, I do not commit myself, and do not intend to suggest, that either ordinary people or ideal rational moral agents deduce their beliefs regarding the betterness relations that hold between different pairs of acts from their beliefs over the set of all possible rankings of all the possible acts available to them.

As Leonard Savage (1972) does, we can define each element of D as the constant act (i.e., an act that gives the same outcome in every state) whose value is this element and require that E includes all the possible—constant and not constant—acts. With this we can treat the agent’s beliefs regarding the betterness relation between constant acts as his beliefs about the betterness relation between outcomes, and the agent’s preferences over constant acts as his preferences over outcomes. For convenience we will use the notation $q(A>B)$ to refer to $q(a_A>a_B)$ when a_A is the constant act that gives A and a_B is the constant act that gives B .

In the interpretation, p represents the agent’s degrees of belief about factual matters in the world, while q represents the agent’s degrees of belief about the betterness relation between different acts. Now, we have the conceptual resources to formally express both the LBC (informally introduced above) and another constraint, the EBC.

Likelihood of Betterness Constraint (LBC):

1. $q(a_i > a_j) > q(a_j > a_i)$ iff $a_i > a_j$ and
2. $q(a_i > a_j) = q(a_j > a_i)$ iff $a_i = a_j$.

Expectation of Betterness Constraint (EBC):

For every two acts, a_i and a_j ,

$$q(a_i > a_j) = \frac{\sum_{\omega_k: a_i(\omega_k) \neq a_j(\omega_k)} p(\omega_k) q(a_i(\omega_k) > a_j(\omega_k))}{\sum_{\omega_k: a_i(\omega_k) \neq a_j(\omega_k)} p(\omega_k)}$$

The agents might form their beliefs in such a way (although I find it psychologically implausible and normatively unappealing), but nothing in the model requires them to do so. This is because I do not assume anything about conditional probabilities; that is the probability of one act being better than another conditional on other betterness relations holding between other acts. Thus, I do not use any information that one gains from access to a specific probability distribution over the set of all possible rankings of the acts and that one does not have if one only has access to the probability of one act being better than another, for all pairs of acts.

Now we can see that the distinction made in the previous section between three types of moral uncertainty on a conceptual level—that is, the distinction between 1) moral uncertainty that can be reduced to uncertainty about non-normative propositions, 2) moral uncertainty that can be reduced to uncertainty about which moral theory is the correct one, and 3) “primitive” moral uncertainty—can be represented formally in a straightforward way: The first kind of moral uncertainty happens when there is no uncertainty regarding the agent’s own preferences, the second happens when there is such uncertainty but all the probabilities, including the conditional probabilities are known to the agent, and the third happens when only non-conditional probabilities are known (or in other words, when the probabilities of conjunctions are not known).

Intuitively, the EBC says that one's degree of belief that one act is better than another should be equal to one's expected degree of belief that this act is better than the other, in case one of the two acts is better than the other. In other words it requires that the agent's degree of belief that one act is better than another is equal to the agent's degree of belief that the world is such that this act is better than the other.

It will be useful to demonstrate how the EBC works, using an example. Consider the following table.

Table 3

	$p(\omega_1) = 0.2$	$p(\omega_2) = 0.3$	$p(\omega_3) = 0.4$	$p(\omega_4) = 0.1$
a_i	A	B	C	B
a_j	B	C	A	B
a_A	A	A	A	A
a_B	B	B	B	B
a_C	C	C	C	C

Suppose the agent's degree of belief that outcome A is better than outcome B (that is that act a_A is better than act a_B) is 0.7, that his degree of belief that B is better than C is 0.8 and that his degree of belief that A is better than C is 0.9. What should his degree of belief be that a_i is better than a_j ? According to the EBC it should be $(0.2 \times 0.7 + 0.3 \times 0.8 + 0.4 \times 0.1) / 0.9 = 0.4666$.

Here is how the calculation goes: firstly the agent should check in which states the two acts give the same outcome and ignore these states. In our example this only happens in state ω_4 . Next, the agent should give each of the remaining states a weight which is equal to its probability and add up his weighted degrees of belief that act a_i is better than act a_j .¹² Lastly, he should normalise this sum by dividing it by the sum of the probabilities of all the states he did not rule out in the first stage. This last move is necessary in order for the agent's degrees of belief to be probabilistic.

To see why, given the EBC, there always exists a mixed act such that the agent believes regarding any other act that it is less likely or equally

¹² Notice that here I used the assumption that two acts cannot be equally good. It is easy to see that if we relax this assumption, the EBC will have to be slightly adjusted, but nothing significant will change.

likely better than this mixed act, let us start with the case of only three acts with regard to which the agent has intransitive preferences, if he obeys the LBC. We can do this by using the following example:

An agent has to choose between three acts that can bring about, in different states of the world, three possible outcomes: that all the 100 inhabitants of village A will die, that all 200 inhabitants of village B will die, or that all 400 inhabitants of village C will die. Assume that the agent is absolutely confident that it is better to save more people than fewer people, thus, $q(A>C)=q(A>B)=q(B>C)=1$. However, the choice he has to make is not between sure outcomes, but between the following three acts:

Table 4

	$p(\omega_1) = 4/9$	$p(\omega_2) = 3/9$	$p(\omega_3) = 2/9$
a_i	B	B	B
a_j	A	C	C
a_k	B	A	C

The agent is following the two conditions mentioned above:

1. EBC: for every two acts, a_i and a_j ,

$$q(a_i > a_j) = \frac{\sum_{\omega_k: a_i(\omega_k) \neq a_j(\omega_k)} p(\omega_k) q(a_i(\omega_k) > a_j(\omega_k))}{\sum_{\omega_k: a_i(\omega_k) \neq a_j(\omega_k)} p(\omega_k)}$$

2. LBC: for every two acts $a_i, a_j, a_i \geq a_j$ iff $q(a_i > a_j) \geq q(a_j > a_i)$.

Now, since $p(\omega_2)+p(\omega_3)>p(\omega_1)$, he believes a_i is better than a_j to degree 5/9. Since $p(\omega_1) >p(\omega_2)$, he believes that a_j is better than a_k to degree 4/7, but since $p(\omega_2) >p(\omega_3)$, he also believes that a_k is better than a_i to degree 3/5 and thus he has intransitive preferences.

We are looking now for a mixed strategy, M, over the three acts such that the agent will believe that M is better than or equal to each one them. We can look at this in the following way. When the agent is using a mixed strategy, he adds some uncertainty to the uncertainty he already suffers from: he transforms any world ω_i to which he gives a positive probability into three worlds, the probability of each one of these being the multiplication of the probability of the original world by the probability that the mixed strategy the agent uses gives to one of the original acts. Here is how this is done in our example:

Table 5

	$p(\omega_1)^*$ M(a _i)	$p(\omega_1)^*$ M(a _j)	$p(\omega_1)^*$ M(a _k)	$p(\omega_2)^*$ M(a _i)	$p(\omega_2)^*$ M(a _j)	$p(\omega_2)^*$ M(a _k)	$p(\omega_3)^*$ M(a _i)	$p(\omega_3)^*$ M(a _j)	$p(\omega_3)^*$ M(a _k)
M	B	A	B	B	C	A	B	C	C
a _i	B	B	B	B	B	B	B	B	B
a _j	A	A	A	C	C	C	C	C	C
a _k	B	B	B	A	A	A	C	C	C

Now, M is preferred or equal to a_i only when the agent believes it is more likely or equally likely that M is better than a_i, i.e., when the sum of the degrees of beliefs that the outcomes that M brings about in every possible world in which M and a_i bring about different outcomes, weighted by the probabilities of these worlds, is higher than this sum for a_i, i.e., since we assumed that the agent’s degrees of beliefs regarding the betterness relations among pure outcomes are all equal to 1, when:

$$p(\omega_1)^*M(a_j) + p(\omega_2)^*M(a_k) \geq p(\omega_2)^*M(a_j) + p(\omega_3)^*M(a_j) + p(\omega_3)^*M(a_k)$$

We can do the same for M in relation to a_j and a_k, and we get three inequalities with three variables. Each inequality can be derived from the other two, but we also know that $M(a_i) + M(a_j) + M(a_k) = 1$. It is easy to see that there is a unique solution to this system in which the equality relation holds for all inequalities. For the values in the example, this solution is when $M(a_i) = M(a_j) = M(a_k) = 1/3$, and in the general case:¹³

$$M(a_i) = (2q(a_j > a_k) - 1) / ((2q(a_j > a_k) - 1) + (2q(a_i > a_j) - 1) + (2q(a_k > a_i) - 1))$$

$$M(a_j) = (2q(a_k > a_i) - 1) / ((2q(a_j > a_k) - 1) + (2q(a_i > a_j) - 1) + (2q(a_k > a_i) - 1))$$

$$M(a_k) = (2q(a_i > a_j) - 1) / ((2q(a_j > a_k) - 1) + (2q(a_i > a_j) - 1) + (2q(a_k > a_i) - 1))$$

¹³ By “the general case” I do not mean only that this solution holds for any p(.), but also that it holds for any q(.), and for any finite number of outcomes (over which the three acts are defined). Finding solutions for more than 3 acts is more difficult. I will show, however, that such a solution always exists.

These values also have an intuitive interpretation, which will be discussed in the next section. The story, however, does not end here, as it is easy to see that for every mixed strategy, such as M, there exist two other acts such that the agent has intransitive preferences over M and these two acts. In our example, this can be done in the following way:

Table 6

	$p(\omega_1)^*$ M(a _i)	$p(\omega_1)^*$ M(a _j)	$p(\omega_1)^*$ M(a _k)	$p(\omega_2)^*$ M(a _i)	$p(\omega_2)^*$ M(a _j)	$p(\omega_2)^*$ M(a _k)	$p(\omega_3)^*$ M(a _i)	$p(\omega_3)^*$ M(a _j)	$p(\omega_3)^*$ M(a _k)
M	B	A	B	B	C	A	B	C	C
N	A	A	B	C	C	A	C	C	C
L	B	A	B	A	C	A	C	C	C

The reasons are identical to the reasons for the intransitivity in the original example. However, notice that N and L are not mixed strategies over the three original acts. Given the set of the original acts and every mixed strategy over them, there is a unique mixed strategy that respects the condition that the agent should never choose a strategy when there exists another strategy available to him that he believes is, more likely than not, better. It seems, then, that in this kind of case the only rational choice for the agent is this mixed strategy.

What happens, though, when the set of available strategies contains more acts? For example, what happens if this set contains the three acts from our example, acts N and L, and every mixed strategy over these five acts? Is it still true that there exists a unique mixed strategy, M, over this set, such that there is no strategy in this set that the agent believes is, more likely than not, better?

The answer to the existence question is yes (I will get back to the uniqueness question soon). To see that, we can think of the agent as playing a game against himself in which the payoffs for every combination of strategies are the agent’s degrees of belief that one of these strategies is better than the other: the intuition is that when the agent has to make a choice, my demand from him is that, given what he chooses, there is no other strategy he could have chosen that he believes will be better. So we can think of it in the following way: the agent looks at the strategies available to him and asks himself—for each

one of them—given that I choose this strategy, will there be a better strategy for me to have chosen? If the answer is yes he should not choose that strategy. It is easy to see that this condition holds for the two players in the game only when they play Nash equilibrium strategies.

Now, since the agent plays against himself, the game is symmetric: the strategies and the payoffs for each combination of strategies for the two players are identical. In the same way, since the two players represent the same agent, the equilibrium must be a symmetric one, since the agent can choose only one strategy. So what we have is a two-player symmetric game and every symmetric game has a symmetric Nash equilibrium (see Nash 1951).

To see things more clearly, let us construct such a game, using our original example. Each player has three pure strategies, a_i , a_j and a_k and the payoff every player gets from choosing an act a , while the other agent chooses act b , is just his degree of belief that a is better than b . Since we assume that the agent ignores, in his reasoning, worlds in which the two acts give the same outcome, we can assign a payoff of $\frac{1}{2}$ to every result in which the two players choose the same pure strategy. So here is the game:

Table 7

	a_i	a_j	a_k
a_i	$\frac{1}{2}, \frac{1}{2}$	$q(a_i > a_j), q(a_j > a_i)$	$q(a_i > a_k), q(a_k > a_i)$
a_j	$q(a_j > a_i), q(a_i > a_j)$	$\frac{1}{2}, \frac{1}{2}$	$q(a_j > a_k), q(a_k > a_j)$
a_k	$q(a_k > a_i), q(a_i > a_k)$	$q(a_k > a_j), q(a_j > a_k)$	$\frac{1}{2}, \frac{1}{2}$

Notice that if the agent has transitive preferences, i.e., if $q(a_i > a_j) \geq \frac{1}{2}$, $q(a_j > a_k) \geq \frac{1}{2}$, and $q(a_i > a_k) \geq \frac{1}{2}$, the only Nash equilibrium is that both players play the pure strategy a_i . However, when the agent has intransitive preferences (which is the case we are interested in), i.e., when $q(a_i > a_j) \geq \frac{1}{2}$, $q(a_j > a_k) \geq \frac{1}{2}$, but $q(a_k > a_i) \geq \frac{1}{2}$, there is no pure strategies Nash equilibrium. However, there is a mixed strategies equilibrium and in this case it is unique.

Now, although this kind of equilibrium is not always unique, the following still holds: given the set of all symmetric Nash equilibrium mixed strategies, the agent always has transitive preferences among them. The reason is simple. It is a well-known fact that in a mixed strategies Nash equilibrium, given the strategy played by player 1, player 2 is indifferent between all the mixed strategies available to him which are defined over the set of all rationalizable strategies. In particular, he is indifferent between all mixed strategies which belong to the set of symmetric Nash equilibrium mixed strategies. Thus, the transitivity requirement is satisfied in a trivial way: the agent is indifferent between all the mixed strategies in the relevant set. So what we have now is a choice rule that respects the requirement that the agent should never choose a strategy if he believes there exists another strategy available to him which is better, which sometimes recommends (i.e., whenever the agent has intransitive preferences) the use of a mixed strategy.¹⁴

To conclude, what we have shown is that if an agent respects the LBC and the EBC, then—even if he holds intransitive moral preferences— if the agent is allowed to use lotteries, there always exists a lottery which he believes is more likely than not better than all other definite acts or lotteries. Thus, for such an agent it seems that the only rational choice will be to choose this lottery (or, if this lottery is not unique, one of the lotteries included in the set).

Recall now the discussion in the first section. One of our lessons from that discussion was that any account—like Broome’s account—that recommends lotteries also in cases where the agent is not morally indifferent between the definite acts which are available to him, must be committed to the claim that sometimes the objectively right thing to do is not to maximize expected goodness. By following the account presented here, we can see that the agent (subjectively) ought to choose a lottery exactly in those cases in which he cannot maximize any quantity anyway, i.e., when his preferences are intransitive.

To be more precise, what I am arguing is that whenever the agent does not suffer from moral uncertainty he should simply choose the

¹⁴ Note that this result does not depend on the LBC and the EBC. Many other decision theories that allow for intransitive preferences can serve. For example, if instead of using the degrees of belief in the betterness relations as the payoffs of the game, we use expected regret levels, the situation will be the same. More generally, Peter Fishburn (1984) has proved that whenever intransitive preferences can be represented by an SSB utility function, this will be the case.

objectively right act. However, when the agent does suffer from some moral uncertainty and does not have an access to degrees of moral value, then if he obeys the LBC he might find himself having intransitive moral preferences. This does not mean that he believes the moral betterness relation is intransitive. We can assume that the agent believes it is transitive. However, since all he can rely on are his beliefs about this relation—in the kind of situations I have pointed to—he has no way to avoid intransitivity. Thus, in the cases where the intransitivity arises, it seems that *the only rational thing for him to do is to choose a lottery*.

So, in my account, choosing a lottery is not an irrational thing to do, but rather—whenever it is justified to choose a lottery—the only rational thing to do. It is clear that in this account there is no need to claim that sometimes the objectively right thing to do is not to choose the best act: one can (but not “must”) argue that one ought always to choose the best act, but that when one is uncertain about which act that is, the only rational thing to do is to use a lottery. Is it also the best thing to do? Well, yes and no. No, in the sense that when choosing a lottery the agent knows for sure that there is another act available to him that brings a higher amount of expected goodness (but he does not know which act that is). Yes, in the sense that—given his uncertainty—this is the only rational thing for him to do, and if we accept that one ought to be rational in one’s moral choices (which we should) then choosing the lottery is the only morally justified act (what I have called “the subjectively right” thing to do).

It turns out that this account also has some nice predictions regarding the kinds of lotteries we ought to use. Some of these will be discussed in the next section.

WHICH LOTTERIES ARE JUSTIFIED?

In this section, I will consider some of the predictions of my account regarding when, and which, lotteries are justified. I will aim to show that by accepting my account we gain an explanation for some judgements that, I think, are intuitive. This, in turn, serves as independent evidence in favour of my account.

Case 1

Let us begin with a paradigmatic case: there are three individuals, i, j, and k, all in need of a kidney. There is only one kidney available and the moral evaluator is uncertain regarding who should get the kidney.

His degrees of beliefs are such, though, that he believes it is more likely than not that i should get the kidney rather than j, it is more likely than not that j should get the kidney rather than k and it is more likely than not that k should get the kidney rather than i. In such a case, it is easy to confirm that my account will recommend the following lottery among i, j, and k:¹⁵

$$M(a_i) = (2q(a_j > a_k) - 1) / ((2q(a_j > a_k) - 1) + (2q(a_i > a_j) - 1) + (2q(a_k > a_i) - 1))$$

$$M(a_j) = (2q(a_k > a_i) - 1) / ((2q(a_j > a_k) - 1) + (2q(a_i > a_j) - 1) + (2q(a_k > a_i) - 1))$$

$$M(a_k) = (2q(a_i > a_j) - 1) / ((2q(a_j > a_k) - 1) + (2q(a_i > a_j) - 1) + (2q(a_k > a_i) - 1))$$

In other words, the weight individual i gets in the lottery, that is the chance that he will get the kidney (denoted $M(a_i)$), should be proportional to the moral evaluator's degree of belief that giving the kidney to j is better than giving it to k. Note that another sensitivity does not hold: it is not the case that the chance that individual i gets in the lottery is proportional to the moral evaluator's degree of belief that giving the kidney to i is better than giving it to j. On the face of it, I find this phenomenon surprising.

Now, this is simply a result of the assumptions presented in the previous section. However, here is one way to make this requirement intuitive. The moral evaluator believes that if k does not get the kidney, i should get it (since he believes that giving the kidney to i is, more likely than not, better than giving it to j). The only reason the evaluator thinks i should not get the kidney is that he believes it is more likely than not that it is better to give it to k than to i. Thus, to the extent that the evaluator believes the kidney should not go to k, he should give it to i. The extent that the evaluator believes the kidney should not go to k is his degree of belief that it is better to give the kidney to j than to give it to k. Thus, it makes sense that the evaluator should give the kidney to i with a probability that is proportional to his degree of belief that k should not get it, i.e., his degree of belief that it is better to give the kidney to j than to k.

¹⁵ This is the case when the agent assigns probability 0 to the possibility of every two acts being morally equivalent. Relaxing this assumption does not change anything substantial, but it does make the mathematical expressions a little bit more complex. For the sake of clarity, then, I chose to use the assumption.

Case 2

Let us consider, now, the simplest case, which is also the one most discussed in the literature. This is the case in which there is no moral uncertainty and the moral evaluator is morally indifferent between two possible acts.¹⁶ For example, consider a kidney case in which there is one available kidney and two candidates, identical in every respect that the evaluator takes to be morally relevant. In this case, although my account allows the use of a lottery, it does not make it strictly (subjectively) superior to either one of the two definite acts (i.e., of giving the kidney to one of the candidates). Broome's account (and any other account like it—in the sense discussed in the first section of this article) does make the lottery that gives equal chances to the two candidates (objectively) morally superior to both either of the two possible definite acts or any other lottery.

Although this is surely an advantage of Broome's account, as intuitively the lottery that Broome's account recommends in this case *is* strictly morally superior to any other possible act, it is not a weakness of my account. This is so because, as explained, my account is not a rival to Broome's account. When my account justifies using a lottery, it is because using a lottery is the subjectively right thing to do. When Broome's account justifies using a lottery it is because doing so is the objectively right thing to do. In the case considered here, there is no moral uncertainty involved and thus my account does not apply. It is not inconsistent with my account, however, to accept the claim that a different account—like Broome's—does apply in this case and does justify using a lottery (in the same way, it is not inconsistent with my account to deny that).

It is important to note, however, that there is a price that Broome must pay here. If the fairness consideration adds some moral value to the lottery in case there is no moral uncertainty and the evaluator is morally indifferent between the two candidates, it should do so also in the case where there is no moral uncertainty, but the evaluator is not morally indifferent between the two candidates.

For example, consider yet another kidney case involving only two candidates, but this time the candidates are identical in everything, apart from the fact that one has a slightly higher chance of a successful operation. According to Broome's account, there must be some cases in which a lottery between the two definite acts would be morally superior

¹⁶ See, for example, Diamond 1967.

to the act of giving the kidney to the candidate with a slightly higher chance of a successful operation.

In order to generate a lottery under Broome's account you can reduce the difference in the chances of a successful operation between the two candidates as much as you want. At some point—if Broome's account is not empty—you will reach a difference in chances such that choosing a lottery between the two candidates will become morally preferred to simply giving the kidney to the one with the (slightly) higher chance of success.

However, if you are consistent in your choices, you will always make the same choice. Thus, if you face a similar choice over and over again you will always prefer the lottery to the option of simply giving the kidney to the candidate with the slightly higher chances. But no matter how small the difference is between the two candidates' chances of a successful operation, after making this decision enough times this will result in preferring a policy that generates more loss of life to one that generates less.

The trade-off has now become clear: if one is willing to accept that in the indifference case, a lottery is not strictly morally superior to the definite acts, one can deny that the fairness consideration is strong enough to lead to morally preferring a policy that generates more loss of life to one that generates less. If, on the other hand, one is willing to accept that sometimes a policy that generates more loss of life is morally superior to a policy that generates less, one can argue that in the indifference case, the lottery is strictly morally superior to any other act.

A third option is to retain both the judgement that in the indifference case the lottery is strictly morally superior to any other act and the judgement that when the evaluator is not morally indifferent and there is no moral uncertainty involved, a lottery is never justified. One can do this by limiting (in a somewhat artificial way) Broome's account (or any other account like it) to cases of indifference, or by arguing that, for some reason, Broome's account does not apply to the case we consider. This move is unattractive for obvious reasons, but these reasons are theoretical, not ethical.

Here, I am not going to argue in favour of any one of the three possible positions I have just presented. The point I want to make is that my account is consistent with all three of them and thus—no matter what your position regarding the above trilemma is—it should

not stop you from accepting my account as valid. The reason for that is, again, that my account is silent regarding the question of the objective rightness of lotteries, and this is exactly what is at stake here.

Case 3

Consider again a single kidney case, but this time there are ten people, i, j, k, and 11... 17, waiting for the kidney. Assume that the evaluator, after thinking about the decision for a while and gathering relevant information, summarises his judgements using the following table:

Table 8

Age	Chances of success	Any other relevant consideration
i	k	j
j	i	k
k	j	i
11... 17	11... 17	11... 17

In other words, the evaluator believes that, from the point of view of the age of the candidates, i is more suited to get the kidney than j, j is more suited than k, and k is more suited than any of 11... 17. However, from the point of view of the chances for a successful operation, k is ranked above i, who is ranked above j, who is ranked above 11... 17. Finally, when the evaluator thinks of any other relevant moral consideration he ranks j above k, k above i, and i above 11... 17.

What should the moral evaluator do? One thing he can do is to try to give a relative weight to each one of the categories and, using these weights, derive a combined ordering. If he manages to do this and get a transitive ordering, I believe he should simply give the kidney to the person ranked at the top, which will be, of course, either i, j, or k.

The problem, though, is that this kind of case is exactly the kind in which the agent might become uncertain regarding which act is the best choice but does not have access to the degrees of moral value of the different acts (this is why the moral uncertainty arises in the first place) and so it might happen that—by following the LBC—he will find that he has intransitive preferences among i, j, and k. In such a case, my

account will suggest a lottery, but this lottery will give a positive chance only to i, j, and k and no chance at all to l1... l7.

To see why this is the case, recall the analogy with a game that I used in the previous section to show why there always exists a lottery that is weakly preferred to any other act. It was demonstrated that, when the agent chooses such a lottery, his choice must constitute a Nash equilibrium in the game he plays against himself.

Now, it is well known that a mixed strategies Nash equilibrium must give a positive chance only to rationalisable strategies, i.e., strategies that can survive the process of iterated elimination of dominated strategies. It is clear that giving the kidney to any of l1... l7 is not a rationalisable strategy because it is dominated by giving the kidney to either i, j, or k. Thus, according to my account, if the agent should use a lottery (which might or might not be the case depending on the agent's beliefs) this lottery must give a positive chance only to i, j, and k.

This seems to me very intuitive. Giving a positive chance to all of the candidates reduces the chances of i, j, and k, and this is so even though the evaluator is sure that it would be wrong to give the kidney to anybody but i, j, or k. So my account, again, gives the "right" recommendation in this case, although it was not designed in any way to do that.

CONCLUSION

I have presented an account of why choosing a lottery over a definite act is sometimes the (subjectively) right thing to do. According to this account, one ought always to choose the best act available when one can. When one cannot, one should use a lottery, and this is because using a lottery is the only rational thing to do in such a situation. So my account succeeds in satisfying both the requirement that moral preferences be rational and the requirement that one ought always to choose the best act available. Moreover, I have argued that the lotteries suggested by my account are the right ones.

One can accept the account presented here for the rightness of lotteries and reject other accounts, but one can also accept my account alongside other accounts as different valid justifications for the use of lotteries. One can also take the account presented here not only as an account of the rightness of lotteries, but also as an account of the fairness of lotteries, but one does not have to do so. If one does, then one can think of being fair as trying the best one can to do the right

thing. If one does not, than this is ok too, as long as one still believes one *ought* to try the best one can to do the right thing.

REFERENCES

- Arrow, Kenneth J., Amartya K. Sen, and Kōtarō Suzumura (eds). 2002. *Handbook of social choice and welfare*. Amsterdam: Elsevier.
- Bradley, Richard. 2007. A unified Bayesian decision theory. *Theory and Decision*, 63 (3): 233-263.
- Broome, John. 1984. Uncertainty and fairness. *Economic Journal*, 94 (375): 624-632.
- Broome, John. 1990. Fairness. *Proceedings of the Aristotelian Society*, 91: 87-102.
- Broome, John. 1991. *Weighing goods*. Oxford: Blackwell Publisher.
- Broome, John. 1994. Fairness versus doing the most good. *The Hasting Center Report*, 24 (4): 36-39.
- Davidson, Donald, J. C. C. McKinsey, and Patrick Suppes. 1955. Outlines of a formal theory of value, I. *Philosophy of Science*, 22 (2): 140-160.
- Diamond, Peter A. 1967. Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: comment. *The Journal of Political Economy*, 75 (5): 765-766.
- Fishburn, Peter C. 1984. SSB utility theory: an economic perspective. *Mathematical Social Sciences*, 8 (1): 63-94.
- Glover, Jonathan. 1977. *Causing death and saving lives*. New York: Penguin.
- Harsanyi, John C. 1978. Bayesian decision theory and utilitarian ethics. *American Economic Review*, 68 (2): 223-228.
- Hooker, Brad. 2005. Fairness. *Ethical Theory and Moral Practice*, 8 (4): 329-352.
- Karni, Edi. 1996. Social welfare functions and fairness. *Social Choice and Welfare*, 13 (4): 487-496.
- Laffond, Gilbert, Jean-François Laslier, and Michel Le Breton. 1993. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5 (1): 182-201.
- Lockhart, Ted. 2000. *Moral uncertainty and its consequences*. New York: Oxford University Press.
- Loomes, Graham, and Robert Sugden. 1987. Some implications of a more general form of regret theory. *Journal of Economic Theory*, 41 (2): 270-287.
- Nash, John. 1951. Non-cooperative games. *Annals of Mathematics*, 54 (2): 286-295.
- Rescher, Nicholas. 1969. The allocation of exotic life saving therapy. *Ethics*, 79 (3): 173-186.
- Saunders, Ben. 2009. A defence of weighted lotteries in life saving cases. *Ethical Theory and Moral Practice*, 12 (3): 279-290.
- Savage, Leonard J. 1972. *The foundations of statistics*. New York: Dover Publications.
- Sepielli, Andrew. 2009. What to do when you do not know what to do? In *Oxford studies in metaethics*, vol. 4, ed. Russ Shafer-Landau. New York: Oxford University Press, 5-28.
- Sher, George. 1980. What makes a lottery fair? *Noûs*, 14 (2): 203-216.
- Steele, Katie. 2006. What can we rationally value? Manuscript presented at the *Third Annual Austin-Berkeley-CMU Formal Epistemology Workshop*. http://fitelson.org/few/few_06/steele.pdf (accessed April 2012).

Ittay Nissan-Rozen is post-doctoral fellow at the Edelstein centre for the history and philosophy of science at the Hebrew University in Jerusalem. His main research interests are in the areas of philosophy of economics, philosophy of probability, Bayesian epistemology, and formal ethics. Contact e-mail: <I.nissan@lse.ac.uk>