(Being fair as) Doing the best one can

Abstract: A new justification for the use of lotteries for the allocation of goods in some situations is presented and contrasted with John Broome's classic justification. Using a formal model for moral decision making under moral uncertainty, it is argued that sometimes the only rational thing for a morally motivated agent to do is to use a lottery.

Introduction

Many people share the intuition that, in some choice situations, using a lottery among (some of) the alternative courses of action open 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, I think, John Broome's[1]. Broome's justification for this intuition is based on the idea that what makes 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. Thus Broome's explanation of what makes a lottery right has two parts: first he presents an account for 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 of lotteries. According to my justification a lottery is justified in some situations when an

[1] See Broome 1990, 1991, 1994, for example. Other discussions of the questions include Hooker 2005, Sher 1980, Saunders 2009, Rescher 1969, Glover 1977.

agent suffers from moral uncertainty, i.e. in some situations when an agent is unsure what is the morally right thing to do (however, not in every situation in which this 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, of course, 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". This will be done in the framework of a simple extension of Leonard Savage's model for decision making[2].

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 can, of course, take a further step and argue that what makes a lottery right according to my account is also what makes it fair. Hence, one can argue that being fair is just doing the best one can to do the right thing, but one does not have to take this further step. I think there might be good reasons to take this further step, but I will not argue for it in this paper. Here I only present a justification for the use of lotteries, not an account of fairness.

Is this a problem for my account? Judging from some of Broome's remarks, it seems that Broome would say that it is, but I disagree and will explain exactly why in the last section of this paper. First, however, I will present my account and contrast it with Broome's. This will be done in the following way: in Section 1, I will present Broome's account as well as some background issues that will serve me later; in Section 2, I will critically discuss an assumption

[2] See Savage 1972.

Broome implicitly uses in his account - namely that it is possible to compare the strength of the moral claims of different people – and will use this discussion to introduce the idea of moral uncertainty; in Section 3, I will use the discussions in the two previous sections to present my account for the rightness of lotteries.

In Section 4, I will consider some of the implications of my account for specific cases and will contrast them with the implications of Broome's account for these cases. I will argue that in some cases my account can be used to complement Broome's account and in other cases my account disagrees with Broome's account. I will argue that in the latter cases the recommendations of my account are more intuitive than those of Broome[3] and that this should be taken as evidence for it.

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, p.87). Broome also holds that because of this fact "…there will certainly be some circumstances where it is better to hold a lottery than to choose the best candidate deliberately" (Broome 1990, p.99).

[3] However, even if one accepts Broome's recommendations, one can still accept my account as a partial explanation for the rightness of lotteries, i.e. one can hold the position that some lotteries are morally right for the reasons Broome presents and some are right for the reasons I present (and some may be right for other reasons). This is, though, not my position. I will argue that Broome's recommendations in the cases I will discuss in Section 4 are wrong, and my recommendations are right.

This claim, by itself, 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 demands 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 the same outcome. This is demonstrated in Table 1:

	ω_1	ω_2
l	A	B
a	A	A
b	B	B

Table 1

The SP demands that if the agent prefers Act a to Act b than 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[4] in moral contexts. However, this is not the strategy Broome adopts and he (as well as others)

[4] This is the position adopted, for example, by Diamond 1967, who first introduced this problem.

has presented very convincing arguments against it[5]. Broome's suggestion for dealing with the problem is different. He 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[6]. By following this suggestion, the SP is not violated because it does not apply. This is demonstrated in 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

Table 2

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.

Notice that, by using this argument, Broome could have consistently argued that there are no cases in which the fair act is not the right act (a claim that he explicitly denies) . Act l can be ranked at the top of an agent’s moral preference ordering and no principle of rationality will be violated. However,

[5] Broome 1984, Section 2.
[6] See Karni 1996 for a similar suggestion.

there is a price for such a move. By including in the description of the outcomes some properties of the acts that (may) bring them about, one violates another axiom of Bayesian decision theory (that is Savage's theory as introduced in his 1972 book), which Broome calls "the rectangular field assumption". This assumption requires that an agent's preference ordering be defined over the set of all possible acts that can be constructed by assigning any one of the possible outcomes to any one of the different states of the world. This assumption must be violated in our example if we follow the "redescribing the outcomes" method, as it is obvious that under the new description there is no possible act that brings the outcome "A achieved by a lottery" in every state of the world.

Now, as Broome stresses, it is hard to treat this latter assumption as a genuine principle of rationality. Its role is rather to make the framework rich enough to allow for the representation theorem to hold. So Broome argues that he prefers to sacrifice this assumption in order to save the SP. I agree, but that does not change the fact that if we violate this assumption we cannot get (at least under Savage's framework) a representation theorem, which means that, if one does accept that sometimes a lottery between two acts should be preferred to both of them, one must hold that an agent with such preferences cannot be described as maximizing the expectation of any quantity – call it goodness, positive moral value, moral utility or any other name you like.

Although Broome does not explicitly claim this, it seems that this is his reason for allowing the right act to differ from the best act (i.e. the act that brings the most good), when fairness considerations are involved. Broome holds that "…goodness is actually fully reducible to betterness; there is nothing more to goodness than betterness." (Broome 1999, p.164). If the right act is always the best act, which is sometimes a lottery, then – in the framework of Savage, which is the one Broome adopts - there can be no goodness function that maximizing its expected value gives this betterness ordering[7].

I have lingered on the discussion above as it will serve me later in order to support the account of the fairness of lotteries that I will put forward. However, nothing in the considerations mentioned explains why lotteries are fair. The discussion only concerned the question of whether choosing a lottery over all available definite acts must violate some principle of rationality. As explained, Broome holds that this is not the case, but he also suggests an account of fairness that explains what can make a lottery fairer than definite choices. Here it is.

Broome takes fairness to be a proportional satisfaction of claims of different people. The satisfaction should be proportional to the strength of the claims in the sense that "…equal claims require equal satisfaction, that stronger claims require more satisfaction than weaker ones, and also – very importantly – that

[7] While still being optimistic regarding future developments, Broome believes that "…none of the other frameworks suggested in the literature …quite solves the problem…" (Broome 1999, p.117). However, see Bradley 2007 for a representation theorem that does. In any case it is worth mentioning that the method of "redescribing the outcomes" has other unattractive features beside the violation of "the rectangular field assumption" to which it leads. For a discussion see Steele 2006.

weaker claims require some satisfaction” (Broome 1999, p. 117). Claims (for some good), according to Broome, are reasons of a special kind to give the good to a specific person: they are reasons that constitutes“…*duties owed to the candidate herself…”* (Broome 1999, p.115; Broome’s Italics). I will not discuss this definition here, but rather will take it as given (Broome himself does not discuss it too much either).

It is important, though, to see how, by using this definition of fairness, Broome is able to justify the use of lotteries. When one has to distribute some indivisible good among a group of people who have claims to this good and when there is not enough of the good to satisfy all claims, no possible distribution will be completely fair. However, instead of choosing the distribution that brings the most good (but that might be extremely unfair), one can choose a fairer distribution by giving each one of the individuals some chance to get the good. This is, claims Broome, “…not perfect fairness, but it meets the requirement of fairness to some extent” (Broome 1999, p.119).

Two points must be emphasised with regard to this account. The first is that it will not always be justified to prefer the fairer option to the option that brings the most good. Sometimes the goodness consideration will override the fairness consideration and sometimes the opposite will hold. Thus, as was claimed before, sometimes the fair act is not the right act and sometimes the right act is not the act that brings the most good.

The second point is that, in just the same way that claims should be satisfied in proportion to their strength when the goods are divisible, in the case of an indivisible good, the chance each person gets in the lottery should be proportional to the strength of his claim to the good. Thus, Broome's account allows for lotteries in non-trivial cases too (i.e. not only when one is indifferent with regard to who should get the good).

There are many unclear points in Broome's argument: what exactly makes a reason to give the good to a person into a claim by this person? Why is it that a chance for a good can substituted for the good itself? Why is it that claims require proportional satisfaction in the first place? However, I am not going to attack any of these points. Instead, I am going to focus my attention on an implicit assumption that Broome makes, namely the assumption that we can, somehow, compare the strength of the claims of different individuals. This assumption is implicit in the requirement for a proportional satisfaction of claims, but Broome does not discuss it. I will do so in the next section.

Interpersonal comparisons of strength of claims and moral uncertainty

It will be useful to remind ourselves, first, of a different – and more famous - problem of interpersonal comparisons: interpersonal comparisons of utility. Here is the usual description of the problem: if an agent's preferences respect the axioms of Bayesian decision theory, then it is possible to represent his choices as a maximization of expected utility for a unique probability function and a utility function which is unique only up to affine transformation. (i.e. if,

for some probability function, a utility function u represents the agent's preferences, then this is true for any utility function of the form v = au +b). This means that if we want to measure an agent's utility from different outcomes, using only (but all) information regarding his preferences over uncertain (as well as certain) prospects, we are not allowed to attach any significance to the zero point and to the unit size of the scale we will obtain (in the same way that we are not allowed to attach such a meaning to the zero point or to the size of units when we measure temperature).

It is easy to see that, by choosing different "b"s in the formula "v=au+b", we change the zero point and that, by choosing different "a"s, we change the unit's size. Since, given a set of rational preferences, we can use any "a" and any "b" we like and represent these preferences as a maximization of expected utility, we should not attach – based only on our information regarding the preference ordering - any significance to the "a" and to the "b" we choose.

A straightforward consequence of this observation is that, using only the utility functions we construct on the basis of rational preferences over uncertain prospects, we cannot justifiably compare the utility levels of different agents on the same scale. Now, those who argue that interpersonal comparisons of utilities are meaningless base their argument exactly on this observation:

since we have no method of making interpersonal comparisons of utilities, they are meaningless[8].

The validity of this argument should not concern us here. What is important for our purpose is the structure of it. This is so because it seems that, in the case of interpersonal comparison of claims, the structure is exactly the opposite. As I will demonstrate, it seems that in the case of claims it is tempting to argue that because sometimes we *do know* how to compare the claims of different people, in other cases *we know we cannot*.

Here is an example. Consider three candidates, A, B and C, claiming one kidney. Assume that the moral evaluator is certain that A is more suited to get the kidney than B (e.g. he is younger, has greater chances for a successful operation and is superior to B in any other respect that the evaluator takes to be relevant). The same holds for B when compared to C, and so also for A when compared to C. However, the only reason that C is in a need of the kidney is because he donated his own kidney to A, a couple of months before the moral evaluator faced the decision.

If you like, you can also imagine that, when C donated his kidney to A, they signed a contract that says that whenever there will be a case in which a kidney can be given to one of them, C should get it. Moreover, you can imagine that, when the evaluator asks A what he think the decision should be,

---

[8] One obvious reply to this argument is that we might be able to use other kinds of information than the agents' preferences in order to make these comparisons (e.g. psychological or biological information), and in fact this is the line taken by Harsanyi (for example see his 1955 paper). Originally, however, the argument was made by Robbins (1934), who held an extreme positivist approach that was immune to such a manoeuvre. This should not concern us here, though.

A admits that he believes that in a choice between him and C, C should get the kidney, but, in a choice between him and B, he should get the kidney, and in the same way B claims that A should get the kidney rather then him, but he should get the kidney rather than C, and C agrees that he should get the kidney rather than A, but B should get the kidney rather than him. What should the evaluator do?

Before I suggest an answer to this question, it is important to try and analyze it using Broome's framework. Intuitively, we know how to compare the strength of the claims of each of two individuals[9]. We know that A's claim to the kidney is stronger than B's, B's is stronger than C's, and C's is stronger than A's. However, *because* we know that, we also know that we cannot compare all these claims on a single scale. If we could, we would not form intransitive judgements regarding which candidate has a stronger claim for the kidney.

Now, Broome believes that "…fairness is concerned only with how well each person's claim is satisfied *compared with* how well other people's are satisfied. It is concerned only with relative satisfaction, not absolute satisfaction" (Broom 1999, p.117, Broome's Italics), but in order to make such comparisons we must assume we can measure the satisfaction of claims on the same scale which in turn commits us to the possibility of comparing the strength of different people's claims on one scale. As Broome argues, "Take a case where all the candidates have claims of equal strength. Then fairness

[9] It is worth mentioning that Broome himself takes kidney transplant cases as paradigmatic cases in which claims arise. See Broome 1990, p.99.

requires equality in satisfaction." (Broome1999, p.117), but in the example above it is tempting to claim that there is no sense in which we can compare the strength of the claims of the three candidates on the same scale and so there is no sense in which we can argue that their claims are equal (or not).

Of course, the moral evaluator can always deny that and claim that there is such a way, but this will commit him to accepting that at least one of the three initial intuitive judgments he formed (which are shared also by the candidates themselves) is wrong – but which one?

It is important to stress that the point of this example is not to suggest that sometimes the relation "morally ought to be chosen rather than"[10] is intransitive. I believe it is transitive. The point, rather, is to suggest that the assumption that interpersonal comparisons of the strength of claims are always possible is implausible, and hence the justification for the fairness of lotteries that is based on it is implausible too. However, even if one is committed to the possibility of interpersonal comparisons of strength claims, nothing in what Broome says tells him how to make these comparisons. I think that the natural reaction to the dilemma in the example is to be uncertain regarding what is the morally right thing to do, and that this is true both if you deny the possibility of interpersonal comparisons of the strength of claims and if you accept it.

[10] I do not use the term "betterness relation" in order to be consistent with Broome's terminology. As explained earlier, Broome takes the betterness relation to represent only some moral considerations – not including fairness considerations – but he does believe (as I do) that both the betterness relation and one's overall moral judgments regarding what one ought to choose (what I called the "morally ought to be chosen rather than" relation) must be transitive (but not necessarily identical). I will claim that there is no need for two different relations, since one ought to always choose the best act available, and I will show that, even with this assumption, sometimes the use of lotteries can be justified.

The idea of being uncertain regarding the question of what is the morally right thing to do is the key element that will help me develop the alternative account of the rightness of lotteries that I will present. However, before discussing this idea more seriously, it is important to stress again what led us to it: this was either the impossibility of making interpersonal comparisons of claims or the difficulty of finding out how that should be done. More generally, I think that, in many cases in which we have to compare the relative strength of different kinds of moral considerations, we are drawn to feeling uncertain regarding what is the morally right thing to do in a situation. This can be either because we are not sure of the appropriate relative weight we should give to each one of the considerations involved, or because we do not believe it is even possible to compare the relative importance of each, but still believe we must make a choice.

This last observation is important because it suggests that the cases in which Broome's account will justify a lottery are roughly the same as those that in my account will justify one. My account will justify a lottery only when one is uncertain regarding what is the morally right thing to do and this happens, roughly, when one has to compare the relative importance of different moral considerations. Comparing the strength of the claims of different people usually falls under this characterization.

This is not always the case. There are cases in which my account would recommend a lottery and Broome's would not, and vice versa. Sometimes

both accounts will recommend a lottery, but different kinds of lottery. I will discuss some of these cases later. However, I believe these will be the exceptional cases and most of the time the two accounts will justify the same lotteries.

Before presenting the way in which being uncertain regarding what the morally right thing to do is leads to using a lottery, it is important to say a little more about this idea. The phenomenon – that is the feeling we sometimes have that we are not sure what the right thing to do is – is very familiar and it will be hard to deny that this is the way at least some people describe their attitudes in some cases. However, at least two different arguments might be used to deny the importance of this phenomenon to ethics.

From the Humean side, one might argue that, although this is the way we sometimes describe our attitudes, it is not the right way to describe them. Our moral judgements, the Humean might argue, are not beliefs, but rather an expression of our moral desires. The levels of certainty we attach to different moral judgements are not really degrees of belief, but rather degrees of desire.

However, by adopting such a position, one loses the ability to explain why people sometimes express intransitive moral judgments, as intransitive judgments are consistent with no desirability function. The Humean might argue that in such cases the agent is simply wrong about his real moral preferences (which are ordered according to their level of desirability), but by

adopting this position he makes it possible again to claim that an agent can be uncertain regarding what is the right thing to do. He simply commits himself to a specific interpretation of what it is that makes something the right thing to do, i.e. that this is what the agent morally desires the most.

From the non-Humean side, one may argue that, although we are sometimes uncertain regarding what is the morally right thing to do, this uncertainty can be reduced to a different kind of uncertainty, i.e. uncertainty regarding which moral theory, or general moral claim, is the correct one. This is the line of thinking adopted by most philosophers in relation to what is now known as "moral uncertainty"[11]. There are two different ways to interpret this position.

The first interpretation is a descriptive one, i.e. it takes the argument to be that when we do feel uncertain regarding what is the morally right thing to do this is always because (and only because) we are uncertain regarding the validity of some general moral principle. I find this interpretation implausible. People can feel uncertain regarding what is the morally right thing to do even if they do not formulate to themselves any general moral principle.

Moreover, people come to believe in moral theories on the basis of their choice recommendations for specific cases. When we find out that a general moral principle or moral theory we accept leads to an unintuitive choice recommendation in a specific case, this is a reason for us to reject this moral principle in its conclusive form. However, the position according to which,

[11] See Lockhart 2000, for example.

whenever an agent feels uncertain regarding what is the morally right thing to do in some situation, it is only because he is uncertain regarding which moral theory or general moral claim is the right one, does not allow for such a reasoning process to take place, as according to this our judgements regarding what we ought to do in specific choice situations are derived from our judgements regarding which moral theory is the right one, and not vice versa.

One might, however, deny this claim on a descriptive level, but accept it as a normative principle: whenever one finds oneself uncertain regarding the right thing to do in a specific situation, one must try and reduce this uncertainty to an uncertainty regarding which one of several moral principles or theories is the correct one.

Such a normative requirement is, however, both vague and destructive for our ethical methodology. It is vague because it is not clear in what sense a moral principle (or a moral theory) is more than the sum of all of its choice recommendations for specific cases. To accept a moral principle, one might argue, is simply to accept all the moral choices it leads to.

It is destructive to our ethical methodology for the same reason that it is implausible as a descriptive account of the way people do their moral reasoning: it forbids us from using our judgements regarding the right thing to do in specific cases as reasons for accepting or rejecting different moral theories. In other words, it denies not only any version of the reflective

equilibrium method, but more generally any form of moral reasoning that aims at moral theories that can have motivational power: without being sensitive – in some way or another – to our moral judgements regarding what we ought to do in specific cases, it is hard to see how a moral theory can be motivational.

So I think the "reducing moral uncertainty to uncertainty about moral theories" claim is implausible at both descriptive and normative levels, and we must consider cases in which agents suffer from uncertainty regarding the morally right thing to do in a way that cannot be reduced to uncertainty regarding the right moral theory. Allowing for such a possibility, however, leads to a problem that will be presented briefly in the next section. This problem is interesting for many reasons, but in this paper I am only going to discuss one positive implication of it: that it leads to some novel predictions regarding the fairness of lotteries.

Moral uncertainty and lotteries

When an agent is uncertain regarding the morally right thing to do in a specific situation, 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, etc. But when she is done with this process, when she has used any source of moral information available to her, then if she is still uncertain regarding what

is the morally right 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 b is better than a, she should choose a over b.

This is not to deny, of course, that an agent can be almost certain that some moral claim, C1, is correct but still (rationally) violate it in her choices because it conflicts with another moral claim, C2, that she thinks is very implausible, but still believes might be true - but if this is true it assigns to a specific act, recommended by C1, a very high level of negative moral value. I think sometimes we do use this kind of “expected level of moral value” calculation and that - when we can - we should do so.

The problem is that in most cases our ability to assign exact levels of goodness or badness to specific acts or outcomes is very limited (even when we do this conditionally on some moral theory being the right one – or even more so in these cases). 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). In such cases, I would argue, after we have used all the information we can regarding the levels of goodness or badness of each possible act, we might take from the point of view of each possible

moral theory we can think of, then if we still feel uncertain regarding what we ought to choose, we must go with the judgement we are most certain about.

Notice that if one accepts this demand, but still wants one's moral choices to be transitive, one commits oneself to the demand that for any three alternatives, A, B and C, if one believes it is more likely than not that A is morally 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 not that A is better than C. This demand also makes intuitive sense, because, if one accepts that the "morally better than" relation is transitive, one's beliefs regarding this relation between A and B and B and C should constrain one's belief about the relation between A and C.

A natural question to ask at this point is how the agent's degrees of belief regarding which alternative is morally better should be affected by his uncertainty regarding the consequences of his acts. In other words, how should uncertainty regarding the world interact with moral uncertainty?

I think a natural answer to this question is that a (rational) agent's degree of belief that one act with uncertain outcomes is better than another such act should be equal to his expected degree of belief that this act is better than the other. In other words, the agent's degree of belief that one act is better than another should be equal to his degree of belief that the world is such that choosing the first act is better than choosing the other.

However, by adding this demand to the two previous ones (i.e. that the agent's moral choices are transitive and that whenever he suffers from moral uncertainty he chooses the alternative he is more certain is the morally right one), we are led to a triviality result.

Here is the basic model (which is basically a simple extension of Savage's 1976 model for decision making). Let $\Omega = \{\omega_1 \ldots \omega_n\}$ be a finite set of possible states of the world. Let p be a probability distribution over Ω. Let D = {A, B, C…} be a set of outcomes and let $A = \{a_{1\ldots}\ a_1\}$ be a set of acts, where an act is a function from Ω to D, and let $\geq^*$ be a regular preference ordering over A (i.e. a complete, reflexive and transitive relation). In addition let $>^{**}$denote the betterness relation between pairs of acts i.e. $>^{**}$ is a binary relation over elements of A such that for any two elements, a_i and a_j, $a_i >^{**} a_j$ or $a_j >^{**} a_i$ [12].

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 simply use the notation > (we will, of course, use the same notation to denote the arithmetical "greater than" relation, but it will be easy enough not to get confused between the two different uses). 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 $a_i >^{**}$ $a_{j.}$

[12] This ignores the possibility that the agent gives a positive probability to the possibility that two acts are equally good or desirable, i.e. that neither of them is better than the other. This assumption makes the discussion simpler and nothing really depends on it.

As Savage does, we can define each element of D as a constant act (i.e. an act that gives the same outcome in every state of the world) whose value is this element, and demand that A include all the possible constant acts. Then 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 $q(a_i > a_j)$ when a_i is the constant act that gives A, and a_j is the constant act that gives B.

Notice that the demand that the agent's moral choice be transitive is already expressed in the model by the demand that the relation $\geq^*$ is transitive. Using the model, we can now formulate the two other demands in the following way[13]:

1. **EBC** (expectation of betterness constraint): For every two acts, a_i and a_j,

$$q(a_i>a_j) = \sum_{wk:ai(wk)\neq aj(wk)} p(w_k)q(a_i(w_k)>a_j(w_k)) / \sum_{wk:ai(wk)\neq aj(wk)} .$$

2. **LBC** (likelihood of betterness constraint)**:** For every two acts a_i, a_j, $a_i \geq^* a_j$ iff $q(a_i>a_j) \geq q(a_j>a_i)$.

[13] Here ai(wk) refers to the consequence act ai brings in the state wk. Notice that the two axioms are consistent with expected utility theory. For an explanation of why this is the case, of why these two axioms express the two conditions presented above, for further justifications of them and for the proofs, see Nissan (under review). To avoid possible confusion, it is worth mentioning here, though, that LBC does not violate expected utility theory since it refers to an element that does not exist in traditional expected utility theory, i.e. to the agent's degree of beliefs regarding the betterness relations.

The triviality result states that these two axioms are satisfied in the model if and only if the following equation holds for every three outcomes, A, B and C, such that A is preferred to B and B is preferred to C:

$$q(A>C) = q(A>B) + q(B>C) - ½ .$$

This means that the agent can be equally certain about his three moral judgements if and only if he is indifferent between all three outcomes. Specifically, he cannot be completely sure about these three moral judgments. Even worse, it is easy to show that, as the number of outcomes increases, the agent's degrees of confidence in each of his moral judgements must decrease more and more and at the limit he must be indifferent between all outcomes except two (the one he ranks at the top and the one he ranks at the bottom, regarding which he should be completely sure that the former is better than the latter).

There are many different ways to interpret this result. Here I will only discuss one. It can be argued that the result shows that when an agent is uncertain regarding what is the morally right thing to do then – except in trivial cases – he must have intransitive moral preferences. What should the agent do in such cases?

Here is one possible answer: if we allow the agent to use mixed strategies, i.e. if we demand that the set of acts available to the agent, A, is convex, that is it must include all the mixed strategies over this set, then 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 demand from the agent to choose such an act. To see this, let us start with the case of only 3 acts with regard to which the agent has intransitive preferences. 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:

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

Table 3

The agent is following the two conditions mentioned, i.e.

1. **EBC**: For every two acts, a_i and a_j,

$$q(a_i>a_j) = \sum_{wk:ai(wk)\neq aj(wk)} p(w_k)q(a_i(w_k)>a_j(w_k)) / \sum_{wk:ai(wk)\neq aj(wk)} .$$

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.

Of course, realizing that, the agent may choose to revise some of his degrees of beliefs so that his preferences will become transitive. If he succeeds in doing this, than this agent can be described as an expected moral value maximizer, i.e. by changing his degrees of beliefs in such a way that his preferences will become transitive and the two axioms will be satisfied, the agent implicitly assigns degrees of moral value to the three possible outcomes. However, the triviality result shows that, as he considers more and more possible outcomes (hypothetical or real), using such a strategy must lead him – at the limit – to be morally indifferent among all acts except two. If we want to avoid this conclusion we must accept that in some situations the agent does have intransitive moral preferences. So, for convenience, we can assume that the agent in our example has already made all the revisions of his degrees of beliefs that he was willing to make.

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 of 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:

	$p(\omega_1)$* $M(a_i)$	$p(\omega_1)$* $M(a_j)$	$p(\omega_1)$* $M(a_k)$	$p(\omega_2)$* $M(a_i)$	$p(\omega_2)$* $M(a_j)$	$p(\omega_2)$* $M(a_k)$	$p(\omega_3)$* $M(a_i)$	$p(\omega_3)$* $M(a_j)$	$p(\omega_3)$* $M(a_k)$
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

Table 4

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 in every possible world in which M and a_i bring 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. Every inequality in this system 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:

$$M(a_i) = ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) / ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) + ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) + ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)})$$

$$M(a_j) = ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)}) / ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) + ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) + ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)})$$

$$M(a_k) = ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) / ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) + ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) + ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)})$$

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 M and these two acts. In our example, for instance, this can be done in the following way:

	$p(\omega_1)$* $M(a_i)$	$p(\omega_1)$* $M(a_j)$	$p(\omega_1)$* $M(a_k)$	$p(\omega_2)$* $M(a_i)$	$p(\omega_2)$* $M(a_j)$	$p(\omega_2)$* $M(a_k)$	$p(\omega_3)$* $M(a_i)$	$p(\omega_3)$* $M(a_j)$	$p(\omega_3)$* $M(a_k)$
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

Table 5

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 to be better than the strategy he actually chose. 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 5 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 it is more likely than not that it is better than M?

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 choose? If the answer is yes he should not choose this 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 2-player symmetric game and every symmetric game has a symmetric Nash equilibrium (See Nash's original 1951 paper).

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 ½ to every result in which the two players choose the same pure strategy. So here is the game:

	a_i	a_j	a_k
a_i	½ , ½	$q(a_i > a_j)$, $q(a_j > a_i)$	$q(a_i > a_k)$, $q(a_k > a_i)$
a_j	$q(a_j > a_i)$, $q(a_i > a_j)$	½ , ½	$q(a_j > a_k)$, $q(a_k > a_j)$
a_k	$q(a_k > a_i)$, $q(a_i > a_k)$	$q(a_k > a_j)$, $q(a_j > a_k)$	½ , ½

Table 6

Notice that if the agent has transitive preferences, i.e. if $q(a_i > a_j) \geq ½$, $q(a_j > a_k) \geq ½$ and $q(a_i > a_k) \geq ½$, 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 $q(a_i > a_j) \geq ½$, $q(a_j > a_k) \geq ½$ but $q(a_k > a_i) \geq ½$, there is no pure strategies Nash equilibrium. However, there is mixed strategies equilibrium and in this case it is unique.

Now, if it is the case that either: 1. for any number of pure strategies in a symmetric 2-players game the mixed strategy symmetric Nash equilibrium is unique, or 2. if it is not unique, given the set of all such mixed strategies, the agent always has transitive preferences among them, then we have a choice rule that respects the demand that the agent should never choose a strategy

such that he believes there exists another strategy available to him which is better, which can always give us a choice and – moreover –sometimes recommends (i.e. whenever the agent has intransitive preferences) the use of a mixed strategy.

It turns out that, although 1 is false, 2 is true. Specifically, the agent must believe regarding every two mixed strategies that are played in a symmetric equilibrium that it is equally likely that either is better than the other, and so must be indifferent between them. The reason is simple. Since any mixed strategy in an equilibrium is a best response against any other strategy, in particular a mixed strategy in a symmetric equilibrium is a best response to a mixed strategy in another symmetric equilibrium (to see that there might be more than one symmetric equilibrium, think of a cycle of 4 alternatives with equal strength of beliefs in each of the betterness relations: both a mixed strategy that gives to two pure strategies probability 1/2 and 1/2 and a mixed strategy that gives each pure strategy 1/4 will satisfy the condition)[14].

To conclude, what we have shown is that if an agent respects the two conditions presented above then, although – except in trivial cases – he must have intransitive moral preferences, if the agent is allowed to use lotteries, there always exist a lottery regarding which he believes no other definite act

[14] Note that this result does not depend on the two conditions presented. 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.

or lottery is more likely than not better than it. Thus, for such an agent it seems that the only rational choice will be to choose one of these lotteries.

Recall that the main challenge for any account – like Broome's - that tries to justify the use of lotteries on grounds of fairness is to deal with the apparent violation of the Sure Thing Principle. Recall also that the only way to reconcile the SP and the claim that it is sometimes morally better to choose a lottery over a definite act was to re-describe the outcomes in such a way that the fairness of the procedure would be incorporated into these. Choosing this strategy, however, makes it impossible – at least in Savage's framework - to describe an agent who prefers a lottery to a definite act as maximizing the expectation of any quantity: goodness, moral utility or what have you.

By following the account presented here, we can see that the agent ought to choose a lottery exactly in those cases when he cannot anyway maximize any quantity, i.e. when his preferences are intransitive. To be more precise, what I am arguing is that whenever an agent has transitive moral preferences he should simply choose the best strategy available to him. However, when the agent suffers from some moral uncertainty, if he obeys the two conditions presented above, he must have intransitive preferences over some acts. This does not mean that he believes the moral betterness relation is intransitive. I have assumed that the agent believes it is transitive. However, since all he can rely on are his beliefs about this relation, he has no way - if he respects the two conditions – to avoid intransitivity. However, in the cases when 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 only 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 right thing to do is not to choose the best act: one ought always to choose the best act, but, when one is uncertain which act this 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 by 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 this 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.

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, when not, and which, lotteries are justified. This is not my argument for my account. The argument had already been presented. However, I am

aware of the fact that some of the steps I took in presenting my account can be rejected. Now, what I want to do is to give you a reason to think twice before doing that. The reason is that by accepting my account we get an explanation for some judgments that I think are intuitive.

Of course, different people have different intuitions and this is particularly true with regard to the case of the moral value of lotteries where our intuitions seem not to be very strong. Thus, my discussion in this section will have to rely heavily on ‘intuition pumps’. I will try to ‘pump’ to you the intuition that the recommendations of my account are correct in the cases I will present. I hope, however, that these will not be misleading intuition pumps, but rather constructive ones[15], i.e. they will push forward intuitions that we have an independent reason to accept and not ones which will lead us to more trouble.

It was argued in the first section that the account presented here for the rightness of lotteries will give – in most cases – similar recommendations to those of Broome’s account. In the second section the reason for that was explained: in my account the use of lotteries will be justified only when moral uncertainty arises and – roughly speaking – cases of moral uncertainty arise when the need to compare the relative strength of different moral considerations arises. Thus, since interpersonal comparisons of the strength of claims are a case like that, both my account and Broome’s may recommend a lottery in cases that fall into this category.

[15] See Dennett 1984 and 1994 for a discussion of the difference between the two.

Now, we have the tools to demonstrate this claim in a more precise way. Consider a case of 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 (as was explained in the last section, these cases are in a sense inescapable). As was also shown in the previous section, in such a case my account will recommend the following lottery among i, j and k:

$$M(a_i) = ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) / ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) + ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) + ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)})$$

$$M(a_j) = ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)}) / ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) + ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) + ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)})$$

$$M(a_k) = ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) / ((2q(a_j > a_k) - 1) * \sum_{wk:aj(wk)\neq ak(wk)}) + ((2q(a_i > a_j) - 1) * \sum_{wk:ai(wk)\neq aj(wk)}) + ((2q(a_k > a_i) - 1) * \sum_{wk:ak(wk)\neq ai(wk)})$$

In other words, the weight individual i gets in the lottery (that is the chances he will get the kidney) should be proportional to the moral evaluator's degree of belief that giving the kidney to j is better than giving it to k. This is simply a result of the model and assumptions presented in the previous section. However, here is one way to make this demand intuitive. The moral evaluator

believes that if k does not get the kidney i should get it (since he believes 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 the evaluator believes the kidney should not go to k, he should give it to i. The extent 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 make 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.

What would Broome's account recommend in this case? According to Broome's account each person should get a chance which is proportional to the strength of his claim for the kidney[16]. What is the strength of the claims of each one of the individuals for the kidney? Nothing in Broome's account tells us how to calculate this, but, given the evaluator's beliefs regarding the betterness relation that holds between the three definite acts open to him, we can argue that i has a claim for the kidney only on the grounds that the evaluator himself believes that it is more likely than not that it is better to give the kidney to j rather than k. This is so since i should get the kidney only to the extent that k does not get it (since the evaluator believes it is more likely than not that k should get it rather than i). Thus, the strength of i's claim for the kidney should be proportional to the evaluator's degree of belief that it is

[16] And note that Broome takes kidney transplant cases to be the most likely candidates for the use of lotteries: "Consider, for instance, life-saving medical treatment such as kidney replacement. It seems plausible that, in these matters of life and death, fairness is particularly important. And it seems plausible that everyone has a claim to life, even if on other grounds some are much better candidates than others" (Broome 1990, p.99).

better to give the kidney to j rather than to k. The same argument holds, of course, for j and k.

Broome does not supply us with a clear criterion for when a reason to give an indivisible good to a person constitutes a claim by this person. However, by accepting my account, a natural criterion arises: in the absence of any reason not to give the indivisible good to the person, any reason to give the good to the person constitutes a claim by this person. When the moral evaluator has transitive moral judgements, the only person who has a claim for the good is, thus, the one that the agent judges to be the best candidate. However, when the moral evaluator has intransitive judgements, each one of the individuals has a claim to the kidney and so – according to Broome's own account – each one should get a chance to get the kidney which is proportional to the strength of his claim.

Notice that, if my account is adopted, the claims discourse is not needed. The only thing that we have to assume in order to support the lottery is that the moral evaluator is rational in his moral choices (that is, rational in the sense presented and defended in the previous section). However, the predictions of my account are that the lottery chosen will be the one Broome's account (under a plausible interpretation) recommends. This is, I think, evidence for my account.

However, as mentioned before, the two accounts will not always give the same recommendations. Here is a case in which they can differ. Consider

again a single kidney case, but this time there are ten people, i, j, k and l1..l7 waiting for the kidney. Assume that the evaluator – after thinking about the decision for a while and gathering relevant information – summarizes his judgments using the following table:

Age	Chances of success	Any other relevant consideration
I	K	j
J	I	k
K	J	i
l1…l7	l1…l7	l1…l7

Table 7

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 l1…l7. 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 l1..l7. Finally, when the evaluator thinks of any other relevant moral consideration he ranks j above k, k above i and i above l1…l7.

What should the moral evaluator do? Well, one thing he can do is to try to give a relative weight to each one of the categories and, using these weights, to derive a combined ordering. If he manages to do that 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 and so it might happen that (using the assumptions of the previous section) 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 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 rationalizable strategies, i.e. to strategies that can survive the process of iterated elimination of dominated strategies. It is clear that giving the kidney to each one of l1…l7 is not a rationalizable 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 be and might not be the case depending on the agent's beliefs) this lottery must give a positive chance only to i, j and k.

What will be the recommendation of Broome's account? Well, first it might be that the goodness considerations in this case will override the fairness

considerations and thus no lottery will be recommended. However, if this is not the case and some lottery will be recommended, then this lottery must give a positive chance to each one of the ten candidates since each one of them has – according to Broome – a claim for the kidney. This seems to me extremely unintuitive to me. Giving a positive chance to each one 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 - among whom he is uncertain who should get the kidney.

Broome would argue that this might be justified because although, in terms of goodness, giving the kidney to one of I1…I7 would be a suboptimal choice, I1…I7 have claims to the kidney and these claims must get partial satisfaction. However, when I try to make intuitive sense of this claim, I find myself imagining a potential conversation between the moral evaluator and each one of the candidates. If the evaluator chose i, j or k, he would have a good response to any complaint made by I1…I7 for not choosing them, i.e. that he believes it is better to give the kidney to someone else. Thus, intuitively, I1…I7 do not have a justified claim for the kidney. This is, however, not the case regarding i, j and k, since, if the evaluator chooses i, for example, then k can justifiably complain that he should have chosen him because the evaluator himself believes it is more likely than not that k is more suited than i. Thus, it sounds reasonable, on the face of it, that only i, j and k have a claim for the kidney.

The above was simply in order to show that claiming that I1…I7 have a claim for the kidney is unintuitive. However, I do not believe that even i, j and k have a claim for the kidney, because for any complaint made by one of them, the evaluator has a good response, namely that if he had given the kidney to the complainer than somebody else would have a justified complaint. The right way to analyze the situation, I believe, is to give up the claims discourse altogether and instead demand from the moral evaluator to do what he believes to be the morally best act. If he does not believe – for any act – that it is more likely than not the best act, he must choose a lottery that he believes no other act (or lottery) is more likely than not to be a better than it. As was explained there always exists such a lottery.

Another family of cases in which my account and Broome's will give different recommendations is the family of cases in which there is no moral uncertainty involved and the agent has transitive preferences. In such cases my account will never recommend a lottery. Broome's account, however, will recommend a lottery sometimes (i.e. when the fairness consideration is strong enough to compensate for the loss of expected goodness generated by choosing the lottery). However, when I try to think of cases where a lottery does seem a better choice than a definite act, I always find myself imagining scenarios where, intuitively, some moral uncertainty is involved. When I think of cases where it is clear what is the best thing to do, I can never bring myself to believe that a lottery is the best choice.

For example, think of yet a third kidney case, but this time involving only two candidates, identical in everything beside the fact that one of them has slightly higher chances 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 chances.

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 between the two candidates' chances for 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. Now, ask yourself: do you still find it intuitive that there is any amount of fairness that can be gained by using a lottery which is sufficient to compensate for the loss of life resulting from choosing the policy that recommends using a lottery over the one that does not? I know some people will be willing to accept this. To these people I have no further argument to suggest.

To conclude, I have presented an account of why choosing a lottery over a definite act is sometimes the 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 the demands both that moral preferences should be rational and 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. Most of the time the same lotteries will be recommended by Broome's account, but sometimes the recommendations will differ. In these cases, I have argued, the recommendations of my account are more intuitive than those of Broome's account.

As mentioned in the introduction, Broome argues that "Sometimes a lottery is the fairest way of distributing a good, and my theory explains, better than any other theory I know of, why this is so. That is the main evidence I offer for it." (Broome 1999. p.111). I hope I have convinced you now that my account explains at least equally well why lotteries are sometimes justified. The question now is whether my account also explains why sometimes a lottery is the fairest way of distributing a good. Nothing in what I have said so far supports this claim. On the contrary, the argument presented here makes it possible to drop any use of the concept of fairness and instead to concentrate solely on goodness.

Regarding another account that explains the rightness of lotteries without making use of fairness considerations (an account presented by Glover 1977 and independently by Rescher 1969), Broome writes "However, I do not think

this argument accounts adequately for the value of lotteries. For one thing, it does not explain their *fairness*." (Broome 1990, p.101). This is a charge that can be directed toward my account too, but I do not believe it is a very troubling charge. I will explain why in the next section.

A methodological remark

The claim that what makes a lottery right is that it is the fairest way of distributing a good can be understood in two different ways. Under one interpretation this is a claim about what fairness is. It has no ethical content, only semantic (or pragmatic) content. Under such a reading, what the claim states is that whatever it is that makes lotteries morally superior to any of the definite choices should be called fairness.

I am not sure that I accept this claim (mainly because I am not sure what should be counted as evidence for it), but I do not mind accepting it. In fact, I believe that by accepting it, and my account for the rightness of lotteries, one can get a nice account of fairness also in contexts that do not lead to the use of a lottery, namely that being fair is doing the best one can to do the right thing[17].

[17] 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, p.331). This 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.

However, I am not going to argue here for such an account. Here I am only concerned with what makes the use of a lottery the right thing to do. This leads me to the second interpretation of the claim.

Under the second interpretation, this is not a claim about what fairness is, but rather about the role fairness plays in moral reasoning. Under such a reading, the claim states that fairness – and different approaches can present different accounts of what fairness is – is an independent moral consideration that should get some weight in one's overall moral assessment of what one ought to do in particular situations, and at least sometimes this weight should lead one to prefer a lottery over a definite choice.

As should be clear from my discussion in the previous section, I do reject the claim under this interpretation. However, as mentioned in the introduction, it is perfectly consistent to accept the claim under this interpretation and to accept my explanation for the rightness of lotteries as well. It simply means that there is more than one reason for using a lottery.

Now, notice that Broome's criticism of Glover's and Rescher's approach ("For one thing, it does not explain their *fairness")*, as well as his main argument in favour of his own account of fairness ("Sometimes a lottery is the fairest way of distributing a good, and my theory explains, better than any other theory I know, why this is so. That is the main evidence I offer for it.") is based on a reasoning from the first reading of the claim to the second: since it is the case that whatever it is that makes lotteries right should be called fairness, and

since Broome's account can explain this, while Glover's and Rescher's account cannot, one should accept Broome's ethical claims regarding which lotteries are fair and which are not, and not Glover's and Rescher's claims.

I find this argument dubious. When considering ethics, I believe, moral intuitions should get priority over linguistic intuitions and this is true even when these linguistic intuitions are intuitions regarding ethical concepts. If one judges the choice recommendations given by some ethical accounts to be the right ones, one ought to follow these even if this requires some revision of the way one uses some ethical concepts. In the same way, the success of an ethical account in explaining the way people use some ethical concepts is not a sufficient reason to accept its choice recommendations (especially when these are unintuitive).

I am particularly convinced that this is the right methodological approach in the case of fairness. This is so because our linguistic intuitions regarding fairness are not very firm. We use the concept in different ways in different contexts and different people use it differently. Now, I think it is an important project to try and find an account of fairness that captures all or at least most of these uses, but this is not an ethical project, but rather a pragmatic or semantic project.

My arguments in favour of my account (presented both in Section 3 and in Section 4) are ethical arguments. They are based either on accepting a moral judgement regarding what is the right thing to do in a particular situation, or on

accepting some judgements regarding what one ought to do in the face of moral uncertainty. Thus, I believe that in so far as one is convinced by my arguments, the possible charge against my account on the grounds that it does not explain the fairness of lotteries should not worry one so much.

Conclusion

When a moral agent is unsure which is the morally right act to choose and is unable to resolve this uncertainty by gathering relevant moral information regarding the degrees of goodness of different acts or outcomes, she must choose an act solely on the basis of her beliefs regarding the moral betterness relations between the acts available to her.

I have argued that such cases are quite common and generally arise when the agent has to make a moral choice in an environment in which more than one moral value is at stake.

I have suggested two plausible conditions that such an agent must obey and pointed to a triviality result that states that, by accepting these two conditions, the agent must find herself having intransitive moral judgements, except in trivial cases.

I have argued that when the agent does have intransitive moral judgements that she cannot find a way to revise, it seems that the only rational thing for her to do is to use a mixed strategy such that she believes no other strategy

available to her is – more likely than not – better than the strategy she chooses. I have shown that such a mixed strategy always exists.

A natural interpretation for mixed strategies in this context is to take them to be actual randomization over (some of) the acts available to the agent. Accepting this interpretation opens the way for a surprising justification for the rightness of lotteries, according to which using a lottery is sometimes the right thing to do because – when it is the right thing to do - choosing this lottery is the best one can do. Accepting this justification for the rightness of lotteries also supplies novel predictions regarding which lotteries we intuitively judge to be morally right.

One can accept the account presented here for the rightness of lotteries and reject Broome's, but one can also accept both these 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 doing the best one can. If one does not, than this is ok too, as long as one believes one *ought* to do the best one can.

Bibliography

Bradley, R. (2007), A Unified Bayesian Decision Theory, *Theory and Decision,* 63:3. pp.233-263.

Broome, J. (1984), Uncertainty and Fairness, *Economic Journal, 94*, pp. 624-632.

Broome, J. (1990), Fairness, *Proceedings of the Aristotelian Society,* 91, pp. 87-102.

Broome, J. (1991), *Weighing Goods,* Oxford: Blackwell Publisher.

Broome, J. (1994), Fairness versus Doing the Most Good, *Hasting Center Report 24*, pp.36-39*.*

Broome, J. (1999), *Ethics out of Economics,* Cambridge University Press.

Dennett, D. (1995). Intuition pumps, *The third culture*, New York: Simon & Schuster.

Dennett, D. (1984). Elbow room: The varieties of free will worth wanting. Cambridge, MA: Bradford Books/MIT Press and Oxford University Press.

Diamond, A. P. (1967), Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment, *The Journal of Political Economics,* 75:5, pp. 765-766.

Fishburn, C.P. (1984), SSB Utility Theory: An Economic Perspective, *Mathematical Social Science,* 8:1, pp.63-94.

Glover, J. (1977), *Causing Death and Saving Lives,* Penguin.

Harsanyi, C.J. (1955), Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility, *The Journal of Political Economics,* 63:4.

Hooker, B. (2005), Fairness, *Ethical Theory and Moral Practice,* 8, pp. 329-352.

Karni, E. (1996), Social Welfare Functions and Fairness, *Social Choice and Welfare,* 13:4, pp.487-496.

Lockhart, T. (2000), *Moral Uncertainty and its Consequences,* Oxford University Press.

Nash, J. (1951), Non-Cooperative Games, *Annals of Mathematics,* 54:2, pp.286-295.

Rescher, N. (1969), The Allocation of Exotic Life Saving Therapy, *Ethics,* 79. pp.173-186.

Robbins, L. (1932), *An Essay on the Nature and Significance of Economic Science*, London: Macmillian.

Saunders, B. (2009), A Defence of Weighted Lotteries in Life Saving Cases, *Ethic Theory and Moral Practice,* 12, pp. 279-290.

Savage L, J. (1972), *The Foundations of Statistic*, Dover Publications

Sher, G. (1980), What makes a Lottery Fair?, *Nous,* 14:2, pp. 203-216.

Steele, K. (2006), What Can We Rationally Value? A Lecture given at the *Third Annual Austin-Berkeley-CMU Formal Epistemology Workshop.*