NOÛS 34:4 (2000) 578–603

The Two-Envelope Paradox, Nonstandard Expected Utility, and the Intensionality of Probability

Terry Horgan
University of Memphis

You are given a choice between two envelopes. You are told, reliably, that each envelope has some money in it—some whole number of dollars, say—and that one envelope contains twice as much money as the other. You don't know which has the higher amount and which has the lower. You choose one, but are given the opportunity to switch to the other. Here is an argument that it is rationally preferable to switch: Let x be the quantity of money in your chosen envelope. Then the quantity in the other is either 1/2x or 2x, and these possibilities are equally likely. So the expected utility of switching is 1/2(1/2x) + 1/2(2x) = 1.25x, whereas that for sticking is only x. So it is rationally preferable to switch.

There is clearly something wrong with this argument. For one thing, it is obvious that neither choice is rationally preferable to the other: it's a tossup. For another, if you switched on the basis of this reasoning, then the same argument could immediately be given for switching back; and so on, indefinitely. For another, there is a parallel argument for the rational preferability of sticking, in terms of the quantity y in the other envelope. But the problem is to provide an adequate account of how the argument goes wrong. This is the two-envelope paradox.

Although there is fairly extensive recent literature on this problem, none of it seems to me to get to the real heart of the matter.[1] In my view, the flaw in the paradoxical argument is considerably harder to diagnose than is usually believed, and an adequate diagnosis reveals important morals about both the nature of probability and the foundations of decision theory. I will offer my own account, in such a way that the morals of the paradox will unfold first and then will generate the diagnosis of how it goes wrong. Thereafter I will briefly pursue some theoretical issues for decision theory that arise in light of the paradox's lessons.

1. Preliminaries.

Some initial observations will provide groundwork for the discussion to follow. First, the paradoxical argument is an expected-utility argument. In decision theory, the notion of expected utility is commonly articulated in something like the following way.[2] Let acts $A_1, \dots, A_m$ be open to the agent. Let states $S_1, \dots, S_n$ be mutually exclusive and jointly exhaustive possible states of the world, and let the agent know this. For each act A_i and each state S_j, let the agent know that if A_i were performed and S_j obtained, then the outcome would be O_{ij} and let the agent assign to each outcome O_{ij} a desirability DO_{ij}. These conditions define a *matrix formulation* of a decision problem. If the states are independent of the acts—probabilistically, counterfactually, and causally—then the *expected utility* of each act A_i is this:

$$U(A_i) = \Sigma_j \, pr(S_j) \cdot DO_{ij}$$

I.e., the expected utility of A_i is the weighted sum of the desirablities of the respective possible outcomes of A_i, as weighted by the probabilities of the respective possible states $S_1, \dots, S_n$.

Second, it appears *prima facie* that the conditions characterizing a matrix formulation of a decision problem are satisfied in the two-envelope situation, in such a way that the paradoxical argument results by applying the definition of expected utility to the relevant matrix. The states are characterized in terms of x, the quantity (whatever it is) in the agent's chosen envelope. Letting the chosen envelope be M (for 'mine') and the non-chosen one be O (for 'other'), we have two possible states of nature, two available acts, and outcomes for each act under each state, expressible this way:

	O contains 1/2x	O contains 2x
Stick	Get x	Get x
Switch	Get 1/2x	Get 2x

Matrix 1

Each of the two states of nature evidently has probability 1/2. So, letting the desirability of the respective outcomes be identical to their numerical values, we can plug into our definition of expected utility:

$$\begin{aligned} U(\text{Stick}) &= [pr(\text{O contains } 1/2x) \cdot D(\text{Get } x)] + [pr(\text{O contains } 2x) \cdot D(\text{Get } x)] \\ &= 1/2 \cdot D(\text{Get } x) + 1/2 \cdot D(\text{Get } x) \\ &= 1/2x + 1/2x \\ &= x \end{aligned}$$

$$\begin{aligned} U(\text{Switch}) &= [pr(O \text{ contains } 1/2x) \cdot D(\text{Get } 1/2x)] \\ &\quad + [pr(O \text{ contains } 2x) \cdot D(\text{Get } 2x)] \\ &= 1/2 \cdot D(\text{Get } 1/2x) + 1/2 \cdot D(\text{Get } 2x) \\ &= 1/2 \cdot 1/2x + 1/2 \cdot 2x \\ &= 1/4x + x \\ &= 5/4x \end{aligned}$$

Third, the operative notion of probability, in the paradoxical argument and in decision theory generally, is *epistemic* in the following important sense: it is tied to the agent's total available information. So I will henceforth call it 'epistemic probability'. Although I will remain neutral here about the philosophically important question of the nature of epistemic probability, lessons that will emerge below from the two-envelope paradox will yield some important constraints on an adequate answer to that question.

Fourth, occasionally below it will be useful to illustrate various points by reference to the following special case of the two-envelope decision situation, which I will call the *urn case*. Here we stipulate that the agent knows that the dollar-amounts of money in the two envelopes were determined by randomly choosing a slip of paper from an urn full of such slips; that on each slip of paper in the urn is written an ordered pair of successive numbers from the set {1,2,4,8,16,32}; that there are an equal number of slips in the urn containing each of these ordered pairs; and that the first number on the randomly chosen slip went into the envelope the agent chose and the second went into the other one. Under these conditions, the acts, states, and outcomes are represented by the following matrix:

	Stick	Switch
M contains 1 and O contains 2	Get 1	Get 2
M contains 2 and O contains 1	Get 2	Get 1
M contains 2 and O contains 4	Get 2	Get 4
M contains 4 and O contains 2	Get 4	Get 2
M contains 4 and O contains 8	Get 4	Get 8
M contains 8 and O contains 4	Get 8	Get 4
M contains 8 and O contains 16	Get 8	Get 16
M contains 16 and O contains 8	Get 16	Get 8
M contains 16 and O contains 32	Get 16	Get 32
M contains 32 and O contains 16	Get 32	Get 16

Matrix 2

Since each of the 10 state-specifications in Matrix 2 has epistemic probability 1/10,

$$U(\text{Stick}) = 1/10(1+2+2+4+4+8+8+16+16+32) = 9.3$$
$$U(\text{Switch}) = 1/10(2+1+4+2+8+4+16+8+32+16) = 9.3$$

Fifth, the paradoxical argument is posed in a way that clearly presupposes that the agent's total available information is symmetrical with respect to the quantity in M and the quantity in O. In virtue of this feature of the decision problem, any rationally eligible assignment of probabilities to possible quantities in the two envelopes must meet the following *symmetry condition*. (The variable 'n' ranges over numerical dollar-amounts that might be in M or in O.)

(n)[pr(M contains n & O contains 2n)
= pr(M contains 2n & O contains n)]

Sixth, discussions of the paradoxical argument frequently take it to presuppose an assignment of probabilities to possible quantities in M and O that satisfies the following condition, which I will call the *asymmetrical symmetry condition* (for short, ASC):

(n)[pr(O contains 1/2n & M contains n)
= pr(O contains 2n & M contains n) = 1/2 pr(M contains n)],

equivalently expressible in terms of conditional probabilities this way:

(n)[pr(O contains 1/2n, given that M contains n) = pr(O contains 2n, given that M contains n) = 1/2].[3]

This condition is asymmetrical in one way—in what it says about O in comparison to what it says about M. But in another way it is symmetrical—in what it says about O's containing 1/2n in comparison to what it says about O's containing 2n. Hence the name.

Often the paradoxical argument is dismissed rather quickly, on the grounds (i) that the ASC is not built into the problem, and/or (ii) that no rational, minimally informed, person would have a probability assignment conforming to the ASC, in this decision situation. These claims about the ASC, I take it, are surely correct. For one thing, there is in fact a *minimal* possible dollar-quantity n that M could contain (viz., 1); so the probability that O contains half this quantity should be zero, rather than being $1/2 \cdot$pr(M contains n). Secondly, for sufficiently huge dollar-amounts that might be in M, it surely becomes more probable that O contains half of such a quantity than that M contains twice that quantity. Thirdly, the actual range of possible money values in M and O is surely finite, since there is only a finite amount of money in the world, and thus there are only finitely many quantities of money that could be in M or in O; yet, given the symmetry condition, the ASC could be satisfied only by an *infinite* range of possible quantities in M and O.[4]

But although these critical observations about the ASC are indeed correct, they would be telling against the paradoxical argument only if that argument really does presuppose a probability assignment conforming to the ASC. I will argue that when the argument is properly interpreted, it makes no such assumption.

2. Clarifying the Paradoxical Argument.

There are various ways to interpret the paradoxical argument, associated with various ways of construing the use of the symbol 'x' in the reasoning. A range of objections can be raised to the argument, depending on exactly how it is interpreted. (Most of the objections allege that the probability assignments in the paradoxical argument are mistaken; one objection charges equivocation.) It will prove instructive to consider a series of different interpretations, along with associated objections.[5] The upshot will be an interpretation that is not subject to any of the objections that arise against the others. In my view, this resulting interpretation is the most charitable one—but in any case, it is the most challenging to defeat.

After considering and setting aside an interpretation which construes 'x' in the argument as a variable of quantification, I will then turn to a series of interpretations that all construe 'x' as a singular term. Each of the successive interpretations in this latter series will be an elaboration and refinement of—rather than a flat-out alternative to—its predecessor. The last in the series will be the interpretation I recommend, and will prove immune to the various objections that can be raised against other interpretations.

2.1. First Construal: 'x' as Variable of Quantification.

At first sight it seems natural to construe 'x' as a variable of quantification in the paradoxical argument. On this interpretation, the overall form of reasoning is this:

> Consider any quantity x, and suppose that M contains x. Then pr(O contains 1/2x) = 1/2 and pr(O contains 2x) = 1/2 .... Hence, for any quantity x, if M contains x then U(Stick) = x and U(Switch) = 1.25x.

So interpreted, however, the step in the argument where it is claimed that pr(O contains 1/2x) = 1/2 and pr(O contains 2x) = 1/2 is just mistaken.

Perhaps this step is an inference relying on the asymmetrical symmetry condition. If so, then the inference goes wrong in two ways. First, the ASC is false, as explained in section 1. Second, the inferential step is fallacious anyway. From the supposition that M contains x, plus the claim (an instantiation of the ASC) that

> pr(M contains x and O contains 1/2x) = pr(M contains x and O contains 2x) = 1/2pr(M contains x),

it does *not* follow that pr(O contains 1/2x) = 1/2 and pr(O contains 2x) = 1/2. That inference would require the supposition that pr(M contains x) = 1; but this claim is false, for each quantity x that might be in M.

Perhaps the inferential step in question relies instead on the assumption

$$(n)[\text{M contains n} \supset \text{pr(O contains 1/2n)} = \text{pr(O contains 2n)} = 1/2],$$

which together with the conditions of the decision problem entails

$$(n)[\text{M contains n} \supset \text{pr(M contains n)} = 1].$$

But these assertions are *thoroughly* false. In the urn case, for example, there is no quantity such that the probability of M's containing *that quantity* is 1, nor is there any quantity such that the probability of O's containing 1/2 of that quantity—or of O's containing twice that quantity—is 1/2. On the contrary, no matter which of the possible quantities 1, 2, 4, 8, 16, or 32 happens to be the actual quantity in M, the probability of that quantity being in M is either 1/5 (if the quantity is either 2, 4, 8, or 16) or else 1/10 (if the quantity is either 1 or 32). Likewise, the probability of half of that quantity being in O is either zero (if the quantity is 1) or else 1/5 (if the quantity is either 2, 4, 8, 16, or 32).

So if 'x' is construed as a variable of quantification, then the paradoxical argument is a non-starter.

2.2. Second Construal: 'x' as Singular Term.

The objection lately mentioned becomes inapplicable, however, if we construe 'x' in the argument not as a variable of quantification, but instead as a singular referring term. The idea is this: when the reasoning begins by saying "Let x be the quantity in M," this remark introduces 'x' as a singular term which is thereby stipulated, in context, to go proxy for the referring expression 'the quantity in M'. It is natural enough to use 'x' this way, since the numerical value of the quantity in M is unknown. (Arguably, this kind of stipulative, context-bound, singular-term use of symbols like 'x' and 'y' is ubiquitous in mathematical practice; for, it is very plausible that this is what typically happens in "solve for x" problems.[6]) On this construal, Matrix 1 can be reformulated this way:

	O contains 1/2(the quantity in M)	O contains 2(the quantity in M)
Stick	Get the quantity in M	Get the quantity in M
Switch	Get 1/2(the quantity in M)	Get 2(the quantity in M)

Matrix 3

On this interpretation, the two possible states of nature are indeed exclusive and exhaustive. Also, given the information available in the decision situation, evidently each of these states does have an epistemic probability of 1/2.

But now the following objection arises. The relevant quantities to consider are the *actual* quantities in the two envelopes. Call the lower actual quantity Lois, and the higher actual quantity Heidi. The referring expression 'the quantity in M' is a *nonrigid* designator—an expression that has different referents in different possible situations. This kind of reference-variation occurs in Matrix 3: in the specification of the first state, and in the column below it, 'the quantity in M' refers to Heidi; but in the specification of the second state, and in the column below it, 'the quantity in M' refers to Lois. So the paradoxical argument commits an equivocation, by employing a nonrigid designator whose referent changes from one state-description to another. These remarks apply, mutatis mutandis, to the original argument employing the symbol 'x', since on the present construal 'x' functions as a singular term going proxy for 'the quantity in M'. The argument's expected-utility calculations are therefore bogus, since 'x' lacks a single constant referent throughout the course of the calculations. To obtain appropriate calculations of expected utility for sticking and for switching, consider the following matrix:

	O contains 1/2(Heidi)	O contains 2(Lois)
Stick	Get Heidi	Get Lois
Switch	Get 1/2(Heidi)	Get 2(Lois)

Matrix 4

Since 1/2(Heidi) = Lois and 2(Lois) = Heidi, when we calculate expected utilities from Matrix 4 we obtain:

U(Stick) = 1/2(Heidi) + 1/2(Lois)
U(Switch) = 1/2(Lois) + 1/2(Heidi)

Thus, the expected utilities of the two acts are indeed the same. The paradox rests on a simple equivocation fallacy.

2.3. Third Construal: 'x' as Rigid Singular Term.

The objection lately mentioned becomes inapplicable, however, if we construe 'x' in the original paradoxical argument as a *rigid* designator—an expression that has the same referent across all possible situations—and thus as going proxy for a referring phrase that itself is to be understood as a rigid designator. I would contend that the phrase 'the quantity in M' can operate this way, with implicit contextual use-parameters rendering it rigid (in context)—and that it *does* operate this way in the paradoxical argument when one begins by saying "Let x be the quantity in the agent's chosen envelope M." But in any case, let us be explicit and stipulative about the matter. Let us hereby use the modifier 'actual' as a rigidifying operator, and let us construe 'x' in the argument as a rigid singular term that goes proxy for the rigid referring expression 'the actual quantity in M'.[7] On this construal, Matrix 1 can be reformulated this way:

	O contains 1/2(the actual quantity in M)	O contains 2(the actual quantity in M)
Stick	Get the actual quantity in M	Get the actual quantity in M
Switch	Get 1/2(the actual quantity in M)	Get 2(the actual quantity in M)

Matrix 5

Each of the states described here apparently has epistemic probability 1/2, given the agent's available information. And we can plug into the definition of expected utility to get the paradoxical outcome. So the original argument remains so far undefeated, once we construe 'x' as a rigid singular term.

But now the critic might object as follows. It is simply not the case that the probability that O contains 1/2(the actual quantity in M) is 1/2, or that the probability that O contains 2(the actual quantity in M) is 1/2. For, whatever the actual quantity in M is, the probability that M contains *that quantity* is considerably less than 1; and accordingly, the probability that O contains half of that quantity (or twice that quantity) is considerably less than 1/2. In short, if 'x' functions as a rigid singular term in the paradoxical argument, then the probability assignments in the argument are just mistaken.

2.4. Fourth Construal: 'x' as Rigid Singular Term Employed *De Dicto*.

The relevant probability claims at issue here are these:

(1) The probability that O contains 1/2(the actual quantity in M) is 1/2.
(2) The probability that O contains 2(the actual quantity in M) is 1/2.
(3) The probability that M contains the actual quantity in M is 1.

Now in effect, the objection lately noted gives these three statements a narrow-scope or *de re* reading, under which the referring expression 'the actual quantity in M' falls outside the probability-operator, thus:

(1*) $(\exists n)[n = 1/2(\text{the actual quantity in M}) \ \& \ pr(\text{O contains n}) = 1/2]$
(2*) $(\exists n)[n = 2(\text{the actual quantity in M}) \ \& \ pr(\text{O contains n}) = 1/2]$
(3*) $(\exists n)[n = \text{the actual quantity in M} \ \& \ pr(\text{M contains n}) = 1]$.

But the defender of the paradoxical argument can reply by saying that although statements (1*)–(3*) are indeed false, the relevant interpretation of statements (1)–(3) gives them the following wide-scope or *de dicto* reading:

(1′) pr(O contains 1/2(the actual quantity in M)) = 1/2
(2′) pr(O contains 2(the actual quantity in M)) = 1/2
(3′) pr(M contains the actual quantity in M) = 1

So each of the two possible states of nature described in Matrix 5 does indeed have probability 1/2; the relevant probability assignments are *de dicto*, not *de re*. And the point carries over to the original paradoxical argument, since ‘x’ there goes proxy for the singular referring expression ‘the actual quantity in M’.

But now the critic might raise the following objection. It can be granted that (3′) is true, because its truth condition is just this:

(3#) (n)[pr(M contains n, given that M contains n) = 1]

But the respective the truth conditions for (1′) and (2′), on the other hand, are these:

(1#) (n)[pr(O contains 1/2n, given that M contains n) = 1/2]
(2#) (n)[pr(O contains 2n, given that M contains n) = 1/2]

And the conjunction of (1#) and (2#) is one formulation of the asymmetrical symmetry condition ASC (discussed and criticized in section 1). So the paradoxical argument mistakenly assumes the ASC. The probability assignments in the argument are therefore mistaken.

2.5. Fifth Construal: ‘x’ as Rigid Singular Term, Employed *De Dicto* to Express Disjunctive Epistemic Possibilities.

The clear-headed defender of the paradoxical argument, I submit, should claim that the argument does *not* presuppose the ASC, that statements (1#)–(3#) do *not* give the truth conditions for statements (1′)–(3′), and that the probability assignments in the argument are correct. The objection lately noted constitutes a challenge to explain how (1′)–(3′) are to be understood, and why they are true. Let us turn to that, focusing first on the urn case, and then generalizing from it.

With respect to the urn case, the following list of statements constitutes a fine-grained specification—expressed in terms of the rigid singular term ‘the actual quantity in M’—of the epistemic possibilities concerning the contents of envelopes M and O:

1. The actual quantity in M = 1 & O contains 2
2. The actual quantity in M = 2 & O contains 1
3. The actual quantity in M = 2 & O contains 4
4. The actual quantity in M = 4 & O contains 2
5. The actual quantity in M = 4 & O contains 8

6. The actual quantity in M = 8 & O contains 4
7. The actual quantity in M = 8 & O contains 16
8. The actual quantity in M = 16 & O contains 8
9. The actual quantity in M = 16 & O contains 32
10. The actual quantity in M = 32 & O contains 16

Each statement on this list has epistemic probability 1/10. Hence, since all the statements are probabilistically independent of one another, the disjunction of the five even-numbered statements on the list has probability 1/2, and the disjunction of the five odd-numbered ones also has one half. But the epistemic probability of the statement

O contains 1/2(the actual quantity in M)

is just the epistemic probability of the disjunction of the even-numbered statements on the list, since each even-numbered disjunct specifies one of the epistemically possible ways that this statement could be true. Likewise, the epistemic probability of the statement

O contains 2(the actual quantity in M)

is just the epistemic probability of the disjunction of the *odd*-numbered statements on the list, since each of the odd-numbered statements specifies one of the epistemically possible ways that *this* statement could be true. *Therefore*, in the urn case, the statements

(1′) pr(O contains 1/2(the actual quantity in M)) = 1/2
(2′) pr(O contains 2(the actual quantity in M)) = 1/2

are true. In both (1′) and (2′), the constituent statement within the scope of 'pr' expresses a *coarse-grained epistemic possibility*, a possibility subsuming exactly half of the ten equally probable fine-grained epistemic possibilities corresponding to the statements on the above list.[8] Each of these two coarse-grained epistemic possibilities does indeed have probability 1/2, since each possibility is just the disjunction of half of the ten equally probable fine-grained epistemic possibilities. There is no reliance here on the asymmetrical symmetry condition. That condition is *not* equivalent to, and is not entailed by, the conjunction of statements (1′) and (2′).

These points generalize from our special case of the two-envelope situation, to the situation as described in the original paradoxical argument. Take any rationally eligible probability assignment *A* to pairs of quantities that might be in M and O respectively, and consider a list of statements like the above list, with one statement on the list for each ordered pair of quantitites which, according to *A*, have non-zero probability of being in M and O respectively. For each statement on the list, let the statement's *counterpart* be the statement obtained by switching the two numerals within it, as in

The actual quantity in M = 2 and O contains 1.
The actual quantity in M = 1 and O contains 2.

The counterpart of every even-numbered statement is an odd-numbered statement, and vice versa. Now, since *A* is a rationally eligible probability assignment, *A* will obey the symmetry condition,

(n)[pr(M contains n & O contains 2n) = pr(M contains 2n & O contains n)]

This being so, each pair of counterpart statements will have the same epistemic probability under *A*. So the disjunction of all the even-numbered statements will have epistemic probability 1/2, and so will the disjunction of all the odd-numbered statements. Thus, under any rationally eligible probability assignment, statements (1′) and (2′) will be true. Far from presupposing the asymmetrical symmetry condition, these statements presuppose only that any rationally eligible probability assignment obeys the symmetry condition—which is indisputable.

At this point someone might object that most of the statements on the list will be metaphysically impossible—and that this fact is knowable by the agent.[9] The objection is that there is something wrong with assigning non-zero epistemic probabilities to the members of a set of statements when one knows that most of those statements are metaphysically impossible—and thereby something wrong with our claims about statements (1′) and (2′).

But this objection is not tenable, as the following example should make evident. Suppose that you are told reliably that there are two men, Sam and Dave, one of whom is the other's father—but you are not told who is the father and who is the son. If this is all the information you have about them, then the appropriate assignments of epistemic probability are these:

pr(Sam is Dave's father) = 1/2
pr(Dave is Sam's father) = 1/2.

That is, there are two epistemic possibilities—Sam's being Dave's father, and Dave's being Sam's father—which are equally probable, relative to the available information. Now, given the widely accepted metaphysical principle of the "necessity of origin" for humans, which asserts that a human being has one's parents essentially (i.e., in all metaphysically possible worlds), one of these two epistemic possibilities—whichever one happens to be non-actual—is metaphysically impossible; and the other one is metaphysically necessary. But they both are epistemically possible nonetheless, and indeed are equally probable.

The upshot is that when the reasoning in the paradoxical argument is properly understood, the probability assignments are correct. There is no equivocation, and no reliance on the asymmetrical symmetry assumption. Diagnosing the problem with the argument is more difficult than it is usually thought to be.

3. Morals of the Paradox.

Let us now draw out some morals of the paradox—first concerning the nature of epistemic probability, and then concerning the foundations of decision theory (specifically, concerning the notion of expected utility). These morals will lead to a diagnosis, in section 4, of how the paradoxical argument goes wrong.

3.1. Epistemic Probability as Intensional.

Some important lessons about epistemic probability have already emerged explicitly in the above discussion: for instance, that epistemic-probability ascriptions can have different truth values when interpreted *de dicto* than when interpreted *de re*, and that epistemic probability sometimes involves a range of epistemic possibilities many of which are not metaphysical possibilities. Closely related to these facts, and just below the surface in the preceding discussion, is a moral that should now be made explicit: epistemic probability is *intensional*, in the sense that the sentential contexts created by the epistemic-probability operator do not permit unrestricted substitution *salva veritate* of co-referring singular terms. Consider the urn case, for example, and suppose that (unbeknownst to the agent, of course) the actual quantity in M is 16. Then the first of the following two statements is true and the second is false, even though the second is obtained from the first by substitution of a co-referring singular term:

pr(M contains the actual quantity in M) = 1
pr(M contains 16) = 1.

Likewise, the first of the following two statements is true and the second false, even though the second is obtained from the first by substitution of a co-referring singular term:

pr(O contains 1/2(the actual quantity in M)) = 1/2
pr(O contains 8) = 1/2.

It should not be terribly surprising, upon reflection, that epistemic probability is intensional in the way belief is, since epistemic probability is tied to available information in much the same way as is rational belief.[10]

Numerous difficult philosophical questions now arise about the ontology of epistemic probability, and about the semantics of epistemic-probability ascriptions—questions similar to those that arise for belief and for belief-ascriptions. For instance, what kinds of items are the ones that possess epistemic probabilities? Do epistemic-probability ascriptions apply to such items *simpliciter*, or only "under a description"?

Fortunately, we need not address such issues here. For present purposes, it will suffice to continue a practice already adopted in the above discussion: speaking of *statements* as items that possess epistemic probabilities, and speaking of *epistemic possibilities* as what are expressed by statements occurring within epistemic-probability contexts. Any adequate philosophical account of epistemic probability should be one that cashes out the operative notion of an epistemic possibility (and does so in a way that accommodates the fact that the epistemic possibilities need not all be metaphysical possibilities); and any adequate account should allow statements to have epistemic probabilities, even if only derivatively by virtue of their relation to whatever items the account treats as the fundamental bearers of epistemic probability.[11]

3.2. Quantifying into Probability Contexts: Canonical vs. Noncanonical Substituends.

Among the important philosophical problems that arise in connection with intensional contexts is the question of how to understand quantification into them. Once again, the issues here with respect to epistemic-probability contexts are similar to those for belief-contexts. Although this complex matter cannot be pursued in detail in the present paper, let me make some observations about it that will prove important below.

It is commonly recognized in the philosophical literature that quantification into belief-contexts appears to work in a way that presupposes a restricted range of allowable substituends for quantified variables that occur within such contexts. There will be some class of *canonical* singular referring terms, and only these are the ones that can allowably be instantiated for quantified variables within belief contexts.[12]

We can expect something similar for quantification into epistemic-probability contexts: an implicit restriction on the allowable kinds of substituends for quantified variables, limiting these to certain canonical singular referring terms. For present purposes, we need to distinguish between canonical and noncanonical singular terms that refer to *numerical quantities*. The principal canonical terms here are *numerals*, expressed either in mathematical notation or in natural language. On the other hand, a referring expressing like 'the actual quantity in M' is evidently *not* canonical—not, at any rate, when its specific referent is not known.

Actual usage bears this out. With respect to the two-envelope situation, for instance, the following statements are both true:

(1) $(n)\neg[\text{pr}(M \text{ contains } n) = 1]$
(2) $\text{pr}(M \text{ contains the actual quantity in } M) = 1$

The reason why (2) is compatible with (1) is that the expression ‘the actual quantity in M’ is not a canonical substituend for the numerical variable ‘n’ within the intensional context created by ‘pr’ in (1), so that

(3) ¬[pr(M contains the actual quantity in M) = 1]

is not a permissible substitution-instance of (1).

The general project of distinguishing canonical from noncanonical singular referring expressions, vis-à-vis epistemic-probability contexts, and explaining the basis for the distinction, emerges as important for the foundations of probability. Although this project cannot be pursued at any length here, two key points are especially pertinent for our present purposes. First, a referring expression is canonical, for an agent, only if the expression’s referent is *epistemically determinate* for the agent, given the agent’s total available information—i.e., the agent knows what item the expression refers to. (The expression ‘the actual quantity in M’ fails this test, for an agent who does not know what quantity is in M.) Second, a *statement* to which an epistemic probability is assigned is epistemically determinate, for an agent, only if the statement expresses a *determinate epistemic possibility* for the agent, given the agent’s total available information—i.e., the agent knows which specific state of affairs the statement expresses. (The statement ‘M contains the actual quantity in M’ fails this test, for an agent who does not know what quantity is in M, because there are various different epistemic possibilities concerning the contents of M.)

3.3. Standard vs. Nonstandard Expected Utility.

The intensionality of epistemic probability has direct consequences for the foundations of decision theory, because it affects the notion of expected utility. To begin with, we should distinguish two kinds of state-specifications, in matrix formulations of decision problems. *Canonical* state-specifications are ones in which all constituent referring expressions are canonical vis-à-vis epistemic-probability contexts, whereas *noncanonical* state-specifications have constituent referring expressions that are noncanonical vis-à-vis such contexts. With regard to the two-envelope decision situation, for instance, a state-specification like ‘O contains 8’ counts as canonical, whereas one like ‘O contains 1/2(the actual quantity in M)’ counts as noncanonical.

Second, we should extend the canonical/noncanonical distinction to the specifications, in matrix formulations of decision problems, of the *outcomes* of the various acts under the various states, and of the *desirabilities* of those outcomes. In the two-envelope situation, for instance, an outcome/desirability specification like ‘Get 8’ is canonical, whereas one like ‘Get 1/2(the actual quantity in M)’ is noncanonical. (Recall that we are identifying desirabilities with numerical outcomes, for simplicity.)

Third, we need to recognize that because expected utility involves epistemic probabilities, and because epistemic-probability contexts are intensional, the available acts in a given decision problem can have several different kinds of expected utility. On one hand is *standard* expected utility, calculated by applying the definition of expected utility to a matrix employing canonical specifications of states, outcomes, and desirabilities. On the other hand are various kinds of *nonstandard* expected utility, calculated by applying the definition to matrices involving various kinds of noncanonical specifications.

Take the urn version of the two-envelope problem, for instance, and suppose that (unbeknownst to the agent, of course) M contains 16 and O contains 32. The standard expected utilities, for sticking and for switching, are calculated on the basis of a matrix employing canonical state-specifications, like Matrix 2 (in section 1). Since each of the 10 state-specifications in Matrix 2 has epistemic probability 1/10,

$$U(\text{Stick}) = 1/10(1+2+2+4+4+8+8+16+16+32) = 9.3$$
$$U(\text{Switch}) = 1/10(2+1+4+2+8+4+16+8+32+16) = 9.3$$

On the other hand, one nonstandard kind of expected utility for the acts of sticking and switching, which I will call *x-based* nonstandard utility and I will denote by 'U^x', is calculated by letting 'x' go proxy for 'the actual quantity in M' and then applying the definition of expected utility to a matrix with noncanonical state-specifications formulated in terms of x, viz., Matrix 1 (in section 1). Since each of the two state-specifications in Matrix 1 has epistemic probability 1/2, and since (unbeknownst to the agent) M contains 16,

$$U^x(\text{Stick}) = x = 16$$
$$U^x(\text{Switch}) = 1.25x = 20$$

Another nonstandard kind of expected utility for the acts of sticking and switching, which I will call *y-based* nonstandard utility and I will denote by 'U^y', is calculated by letting 'y' go proxy 'the actual quantity in O' and then applying the definition of expected utility to a matrix with noncanonical state-specifications formulated in terms of y, viz.,

	M contains 1/2y	M contains 2y
Stick	Get 1/2y	Get 2y
Switch	Get y	Get y

Matrix 6

Since each of the two state-specifications in Matrix 6 has epistemic probability 1/2, and since (unbeknownst to the agent) O contains 32,

U^y(Stick) = 1.25y = 40
U^y(Switch) = y = 32

There is nothing contradictory about these various incompatible expected-utility values for sticking and switching in this decision problem, since they involve three different *kinds* of expected utility—the standard kind U, and the two nonstandard kinds U^x and U^y.

Fourth, since a distinction has emerged between standard expected utility and various types of nonstandard expected utility, it now becomes crucial to give a new, more specific, articulation of the basic normative principle in decision theory—the principle of *expected-utility maximization*, prescribing the selection of an action with maximum expected utility. This principle needs to be understood as asserting that rationality requires choosing an action with maximum *standard* expected utility. Properly interpreted, therefore, the expected-utility maximization principle says nothing whatever about the various kinds of nonstandard expected utility that an agent's available acts might also happen to possess.

4. Diagnosis of the Paradox.

We are now in a position to diagnose the problem with the paradoxical argument. Since the kind of expected utility to which the argument appeals is U^x—i.e., x-based nonstandard expected utility—the principal flaw in the argument is its reliance on a mistaken normative assumption, viz., that in the two-envelope decision problem, rationality requires U^x-maximization. Thus, given that U^x is the operative notion of expected utility in the paradoxical argument, the reasoning is actually correct up through the penultimate conclusion that the expected utilities of sticking and switching, respectively, are x and 1.25x. But the mistake is to infer from this that one ought to switch.

Equivocation is surely at work too. Since the unvarnished expression 'the expected utility' is employed throughout, the paradoxical argument trades on the presumption that the kind of expected utility being described is *standard* expected utility. This presumption makes it appear that the normative basis for the final conclusion is just the usual principle that one ought rationally to perform an action with maximal expected utility. But since that principle applies to standard expected utility, whereas the argument is really employing a nonstandard kind, the argument effectively equivocates on the expression 'the expected utility'.

Let me add some further remarks, in order to clarify further both the nature of x-based and y-based nonstandard expected utility in the two-envelope problem, and also the reasons why it is normatively inappropriate (for this problem) to require the maximization of either one. To begin with, it is important to appreciate the way the intensional and the nonintensional interact, within the calculation of U^x and of U^y. On one hand, the states, outcomes, and desirabilities specified in a matrix formulation of a decision problem are what they are, independently of how they are specified; intensionality does not intrude.[13] The epistemic probabilities, on the other hand, attach to the *specifications* of the states (or, if you like, to the states *as specified*); here is where intensionality plays a role.

To illustrate this interplay of the intensional and the nonintensional aspects of expected utility, let us return to the urn version of the two-envelope problem, and suppose once again that M contains 16. Here the outcome for switching specified by 'Get 2x' is identical to the outcome specified by 'Get 32', even though the agent doesn't know what this specific outcome is; likewise, the desirability of this outcome for the agent is 32, even though the agent doesn't know this (since he doesn't know what the outcome is). On the other hand, the probability that attaches to the state-specification 'O contains 2x' is 1/2, whereas the probability that attaches to the state-specification 'O contains 32' is only 1/10.[14]

Because the probability aspect of expected utility is intensional whereas the desirability aspect is not, the following thing can happen. (1) A matrix formulation of a decision problem employs noncanonical specifications of the states, outcomes, and desirabilities—with the items thus specified being epistemically indeterminate for the agent. (2) The set of states specified in the matrix exhibits an *asymmetry* that extends to the outcomes and their desirabilities, and yet does not reflect any corresponding asymmetry in the agent's available information. Nonetheless, (3) each of the state-specifications in the matrix has epistemic probability 1/2.

This happens with Matrix 1. The actual quantity in M is held *fixed* in each state, whereas the quantity in O differs across the two states—even though this asymmetry does not reflect any asymmetry in the available information. Nonetheless each of the states, *as specified*, has epistemic probability 1/2. Because of the asymmetry, $U^x(\text{Switch}) = 5/4 \cdot U^x(\text{Stick})$.

It also happens with Matrix 6. Here the actual quantity in O is held fixed in each state, whereas the quantity in M differs across the two states—once again, an asymmetry in the states that does not reflect any asymmetry in the available information. Nonetheless each of the states, *as specified*, again has epistemic probability 1/2. Because of this asymmetry, $U^y(\text{Stick}) = 5/4 \cdot U^y(\text{Switch})$.

U^x and U^y thus exhibit complementary kinds of asymmetry, in a decision problem in which the available information about the two envelopes is fully symmetrical. Because of this structural feature they both exhibit in offsetting ways, it would be normatively inappropriate to require the maximization of either type of nonstandard expected utility.

5. On Persistent Recalcitrant Intuitions.

Although the diagnosis in section 4 evidently suffices as an explanation of how the paradoxical argument goes wrong, I think there is more to say about its seductive intuitive pull—a pull it exerts, curiously enough, despite being obviously mistaken.

In order to bring to light one likely source of this seductiveness, let us consider the following variant of the situation. After you initially choose an envelope, you are told the quantity in this envelope, and *now* you are given the opportunity to switch. In this scenario, standard expected utility dictates switching. For, if A is the known amount of money in your chosen envelope, then U(Stick) is now just the known dollar-quantity A, and U(Switch) is now 1/2(1/2A) + 1/2(2A) = 1.25A. Hence you ought rationally to switch, since the normative principle of U-maximization requires it.[15]

Notice, however, that in the original two-envelope situation, you already know that this reasoning *would* be correct if you *were* to learn the quantity x in your own envelope. "But since I know *this* already without knowing x," you might think, "surely I thereby know already that it is rationally preferable to switch, even though I do not know the actual value of x." This line of thought has considerable intuitive appeal, despite the simultaneous and conflicting intuition that it is surely mistaken.

But this seductive line of reasoning is indeed mistaken, of course. Prior to being told the amount in your envelope, your information about the two envelopes is thoroughly symmetrical, and hence neither choice is rationally preferable to the other. Being told the quantity in your own envelope would make your information about the envelopes crucially asymmetrical, in a way that would make switching have a higher standard expected utility than sticking. But even though you know in advance that your information *would* become relevantly asymmetrical if you were told the actual quantity in your envelope, this subjunctive knowledge by itself does not create the requisite informational asymmetry. After all, you have matching subjunctive knowledge in the other direction: you know in advance that if you were told the quantity in the *other* envelope, then *sticking* would then acquire a standard expected utility 5/4 that of switching.[16] Actual symmetry in one's information is what matters, counterfactual asymmetry notwithstanding.

6. Nonstandard Expected Utility in Decision Theory.

Now that the distinction has emerged between standard expected utility and various kinds of nonstandard expected utility, important new theoretical questions arise in its wake: Is it sometimes normatively appropriate to require the maximization of certain kinds of nonstandard expected utility? If so, then under what circumstances?

Such questions are of interest for at least two reasons. First, the use of an appropriate kind of nonstandard expected utility might sometimes provide a suitable shortcut-method for deciding on a rationally appropriate action in a given decision situation. A correctly applicable kind of nonstandard expected utility might well employ a much more coarse-grained set of states, thereby simplifying calculation.

Second (and more important), maximizing a certain kind of nonstandard expected utility might sometimes be rationally appropriate in a given decision situation even if the available acts *lack* standard expected utilities—i.e., even if the agent does not possess a standard probability assignment for a suitable set of exclusive and exhaustive states of the world. (By a *standard* assignment of epistemic probabilities, I mean an assignment of probabilities to the states *as canonically specified*.) Indeed, perhaps there are decision situations in which it is rationally appropriate to maximize a certain kind of nonstandard expected utility even when the agent's total available information makes it rationally *inappropriate* to adopt any standard probability assignment, because numerous candidate-assignments all are equally rationally eligible.

6.1. E^z-Maximization in the Two-Envelope Scenario.

The two-envelope situation itself is a case in point. The official description of the situation does not provide enough information to uniquely fix a standard probability-assignment that generates standard expected utilities for sticking and switching. (In this respect, the original problem differs from our special case, the urn version.) Nevertheless, the following is a perfectly sound expected-utility argument for the conclusion that sticking and switching are rationally on a par. Let z be the lower of the two quantities in the envelopes, so that 2z is the higher of the two. Then the epistemic possibilities for states and outcomes are described by the following matrix:

	M contains z and O contains 2z	M contains 2z and O contains z
Stick	Get z	Get 2z
Switch	Get 2z	Get z

Matrix 7

The two state-specifications in Matrix 7 both have probability 1/2. Hence the expected utility of sticking is $1/2z + 1/2(2z) = 3/2z$, whereas the expected utility of switching is $1/2(2z) + 1/2z = 3/2z$. So, since these two acts have the same expected utility, they are rationally on a par.

The soundness of this argument is commonly acknowledged in the literature. What is not commonly acknowledged or noticed, however, is that the notion of expected utility employed here is a nonstandard kind. I will call it *z-based* nonstandard expected utility, and I will denote it by U^z. In order to illustrate the fact that U^z differs from standard expected utility U, return to the urn case, and suppose that (unbeknownst to the agent, of course) the actual lower quantity z in the envelopes is 16 (and hence the actual higher quantity in the envelopes, 2z, is 32). Then, as calculated on the basis of Matrix 2 (in section 1), U(Stick) = U(Switch) = 9.3. However,

$$U^z(\text{Stick}) = 1/2z + 1/2(2z) = 1/2 \cdot 16 + 1/2 \cdot 32 = 24$$
$$U^z(\text{Switch}) = 1/2(2z) + 1/2z = 1/2 \cdot 32 + 1/2 \cdot 16 = 24$$

And with respect to the original two-envelope situation (as opposed to the urn case), there are *no* such quantities as U(Stick) and U(Switch), since there is not any single, uniquely correct, assignment of standard probabilities over canonically-specified epistemic possibilities.

Why is it the case, in the two-envelope situation, that a rationally appropriate action must be a U^z-maximizing action? Because the agent's information concerning the quantities in the two envelopes is entirely symmetric: the agent has no reason to believe that either of the envelopes is more likely than the other one to contain the higher actual value, or the lower actual value. Given this informational symmetry, any rationally eligible assignment of probabilities to canonically-specified potential quantities in the two envelopes must satisfy the symmetry condition,

(n)[pr(M contains n & O contains 2n) = pr(M contains 2n & O contains n)]

Thus, even though different standard probability assignments would yield different pairs of standard expected-utility values for the available acts, the symmetry condition guarantees that for any rationally eligible standard probability assignment *A*, U_A(Stick) = U_A(Switch). (Here, U_A is the standard expected utility as calculated on the basis of *A*.) So, since the two available actions are ranked the same way relative to one another (viz., as rationally on a par) under every rationally eligible standard probability assignment to epistemic possibilities for M and O, and also are ranked that same way by U^z, it is rationally appropriate—in the given decision problem—to employ the normative principle of U^z-maximization. Thus, even though numerous rationally eligible standard probability assignments are all *equally* eligible—so that the available acts do not possess standard expected utilities—rationality requires the selection, in this situation, of a U^z-maximizing act.[17]

6.2. E^x-Maximization in the Coin-Flipping Scenario.

I turn next to some structurally parallel observations about a decision situation somewhat different from the original one, and occasionally discussed in the literature. Suppose you are given an envelope M, and there is another envelope O in front of you. You are reliably informed that M has a dollar-amount of money in it that was chosen by some random process; that thereafter a fair coin was flipped; and that if the coin came up heads then twice the quantity in M was put into O, whereas if the coin came up tails then half the quantity in M was put into O. I will call this the *coin-flipping situation*, in contrast to the *original situation* that generates the two envelope paradox.

In this coin-flipping situation, you ought rationally to switch—as has been correctly observed by those who have discussed it.[18] Moreover, the form of reasoning employed in the original two-envelope paradox not only yields the correct conclusion, but in *this* situation also appears to be a perfectly legitimate way to reason one's way to that conclusion.

This fact is acknowledged in the literature. What is not acknowledged or noticed, however, is that the notion of expected utility employed here is a non-standard kind, viz., U^x. In order to illustrate the fact that U^x differs from standard expected utility in this problem, consider the following coin-flipping variant of the urn case. Suppose that the agent knows that the dollar-amount in his own envelope M was determined by randomly choosing a slip of paper from an urn full of such slips; that on each slip in the urn is written one of the numbers in the set {2,4,8,16,32}; and that there are an equal number of slips in the urn containing each of these numbers. The agent also knows that after the quantity in M was thus determined, the quantity in O was then determined by coin-flip, with twice the quantity in M going into O if the coin turned up heads, and half the quantity in M going into O if the coin turned up tails. Under these conditions, the standard expected utilities are calculated on the basis of the following matrix:

	Stick	Switch
M contains 2 and O contains 1	Get 2	Get 1
M contains 2 and O contains 4	Get 2	Get 4
M contains 4 and O contains 2	Get 4	Get 2
M contains 4 and O contains 8	Get 4	Get 8
M contains 8 and O contains 4	Get 8	Get 4
M contains 8 and O contains 16	Get 8	Get 16
M contains 16 and O contains 8	Get 16	Get 8
M contains 16 and O contains 32	Get 16	Get 32
M contains 32 and O contains 16	Get 32	Get 16
M contains 32 and O contains 64	Get 32	Get 64

Matrix 8

Since the probability is 1/10 for each of the canonically specified states in Matrix 9, the standard expected utilities are

$$U(\text{Stick}) = 1/10(2+2+4+4+8+8+16+16+32+32) = 1/10(124) = 12.4$$
$$U(\text{Switch}) = 1/10(1+4+2+8+4+16+8+32+16+64) = 1/10(155) = 15.5$$

On the other hand, suppose that (unbeknownst to the agent, of course) the actual quantity in M is 32. Then the x-based nonstandard expected utilities are

$$U^x(\text{Stick}) = x = 32$$
$$U^x(\text{Switch}) = 1/2(1/2x) + 1/2(2x) = 1/2\cdot 16 + 1/2\cdot 64 = 40$$

And with respect to the original coin-flipping two-envelope situation (as opposed to our urn version of it), there are *no* such quantities as U(Stick) and

U(Switch), since there is not any single, uniquely correct, assignment of standard probabilities over canonically-specified epistemic possibilities.

Why is it the case, in the coin-flipping version of the two-envelope situation, that a rationally appropriate act must be a U^x-maximizing act? Because in the given decision situation, the agent's information about the quantities in the two envelopes is asymmetric in one respect and symmetric in another. It is asymmetric with respect to how the quantities in the two envelopes were determined: first a quantity was chosen at random to go into the agent's own envelope M, and then either twice or half that quantity went into the other envelope O, depending on the outcome of a fair coin-flip. But the information is symmetric with respect to the possible outcomes of the coin-flip. Given these facts, any rationally permissible assignment of probabilities to canonically-specified potential quantities in the two envelopes must satisfy the asymmetrical symmetry condition,

(n)[pr(O contains 1/2n & M contains n) = pr(O contains 2n & M contains n) = 1/2 pr(M contains n)]

Even though different standard probability assignments would yield different pairs of standard expected-utility values for the available acts, the asymmetrical symmetry condition guarantees that for any rationally eligible standard probability assignment *A*, U_A(Switch) = 1.25 · U_A(Stick).[19] So, since the two available acts stand in the same ratio-scale ranking under every rationally eligible standard probability assignment to the epistemic possibilities for M and O, and also are ranked this same way by U^x, it is rationally appropriate—in the given decision problem—to employ the normative principle of U^x-maximization. Thus, even though numerous rationally eligible standard probability assignments are all *equally* eligible—so that the available acts do not possess standard expected utilities—rationality requires the selection, in this situation, of a U^x-maximizing act.

6.3. Maximizing Nonstandard Expected Utility: A Normative Principle.

In light of sections 6.1 and 6.2, let me now formulate a general normative principle for the maximization of various kinds of nonstandard expected utility in various decision situations. For a given decision problem, let δ be a singular referring expression that is epistemically indeterminate given the total available information, and hence is noncanonical. Let U^δ be a form of nonstandard expected utility, applicable to the available acts in the decision situation, that is calculated on the basis of a matrix employing noncanonical state-specifications formulated in terms of δ. Assume that for the given decision situation, there is at least one rationally eligible standard probability assignment to epistemically possible states of nature. Then rationality requires the maximization of U^δ just in case the following condition obtains:

> There is a unique ratio-scale ordering O of available acts such that (i) for every rationally eligible standard probability assignment *A* to epistemically possible states of nature for the given decision situation, U_A ranks the available acts according to O, and (ii) U^δ also ranks the available acts according to O.[20,21]

This normative principle dictates U^z-maximization in the original two-envelope situation, but not U^x-maximization or U^y-maximization. It dictates U^x-maximization in the coin-flipping version of the two-envelope situation, but not U^y-maximization or U^z-maximization.[22] It has applications not only as an occasional short-cut method for rational decision-making that is simpler than calculating standard expected utility, but also (and much more importantly) as a method for rational decision-making in certain situations where the available acts have no standard expected utilities at all.[23]

Notes

[1]See, for instance, Nalebuff 1989, Cargile 1992, Castell and Batens 1994, Jackson *et al.* 1994, Broome 1995, Arntzenius and McCarthy 1997, Scott and Scott 1997, and Chalmers unpublished.

[2]See, for instance, Jeffrey 1983.

[3]The two formulations are equivalent because the conditional probability of A given B, pr(A/B), is defined this way:

$$pr(A/B) = \frac{pr(A\&B)}{pr(B)}$$

[4]Because it is often assumed in the literature that the paradoxical argument presupposes the ASC, it is also often assumed that the form of reasoning employed in the argument really only has theoretical interest for probability distributions over infinitely many states of nature. Some discussions therefore focus exclusively on whether or not it is mathematically possible for there to be such infinite-range probability distributions that satisfy the ASC and also satisfy the other conditions of the decision problem. Not surprisingly, these discussions can become mathematically quite technical. But mathematical debate about infinite-range probability distributions involving the ASC is really a sideshow as far as the paradox itself is concerned, because in effect this debate already assumes that what's wrong with the paradoxical reasoning in the original two-envelope situation is that it mistakenly presupposes the ASC.

[5]I will not attempt to attribute these various interpretations and objections to specific authors who have written about the paradox, partly because of interpretative issues concerning the authors' texts themselves, and partly because sometimes several different construals of the paradoxical argument may be operating at once within a given text.

[6]I owe this observation to John Tienson.

[7]The expression 'dthat' is employed this way in Kaplan 1978.

[8]Likewise, within statement (3′), the constituent statement within the scope of 'pr'—viz., 'M contains the actual quantity in M'—expresses a coarse-grained epistemic possibility subsuming *all ten* of the fine-grained epistemic possibilities corresponding to the statements on the list.

[9]Since 'the actual quantity in M' is a rigid designator, it designates its actual-world referent at all metaphysically possible worlds. Thus, at least eight of the ten above-listed epistemic possibilities (nine, if the actual quantity in M is either 1 or 32) are not metaphysical possibilities.

[10]This certainly should not be surprising to those who think that epistemic probability is just degree of belief, or rational degree of belief.

[11]If *propositions* are taken to be the fundamental items that possess epistemic probability, then the operative notion of proposition will have to satisfy two constraints. First, propositions must be sufficiently fine-grained that the statements 'M contains the actual quantity in M' and 'M contains 16' express different propositions even if the actual quantity in M is 16. Second, some propositions will correspond to epistemic possibilities that are not metaphysical possibilities. Certain ways of construing a proposition—for instance, as a set of metaphysically possible worlds—evidently violate these constraints.

[12]On this theme see, for instance, Follesdal 1967, Kaplan 1969, and Quine 1969.

[13]At any rate, intensionality does not intrude in any way that is directly relevant to the two-envelope problem. See the next note.

[14]It might be thought that since epistemic probability is intensional, desirability is really intensional too—and in the same way. But the concept of desirability that figures in the notion of x-based nonstandard utility allows substitutivity *salva veritate* in a way that the notion of epistemic probability does not. To see that this is so, notice that the calculation of U^x(Stick) takes for granted that

(1) D(Get the actual quantity in M) = the actual quantity in M

But the second occurrence of the phrase 'the actual quantity in M' in (1) is clearly extensional, since this occurrence does not occur within the scope of either the probability operator 'pr' or the desirability operator 'D'. Hence, from (1) and

(2) the actual quantity in M = 16

it follows that

(3) D(Get the actual quantity in M) = 16

But of course,

(4) D(Get 16) = 16

And from (3) and (4) it follows by the transivity of identity that

(5) D(Get the actual quantity in M) = D(Get 16)

Thus, within contexts created by the desirability operator 'D', canonical referring terms like '16' are intersubstitutable *salva veritate* with co-referring noncanonical referring terms like 'the actual quantity in M'. (Does this show that the operative notion of desirability in U^x, and in decision theory more generally, is *completely* non-intensional? No. But at the very least, it does show that desirability is not intensional *in the same way* as epistemic probability—and in particular, is not intensional in a way that bears directly on the present diagnosis of the two-envelope paradox.)

[15]This argument assumes that in the original two-envelope situation, and thus also in the variant situation just described, (1) if you *were* to learn the actual quantity x in your initially chosen envelope, then you *would* still consider it equally likely that the other envelope contains twice this amount or half this amount; and (2) you know this. Are these suppositions implicit in the usual description of the two-envelope situation? Not clearly so, and nothing in my prior discussion requires it. (Certainly the following two statements are compatible: (i) the x-based state-specifications "O contains 1/2x" and "O contains 2x" each have epistemic probability 1/2; (ii) the actual quantity in your initially chosen envelope is such that if you were to learn what it is, then you would not believe that the other envelope is equally likely to contain either half that amount or twice that

amount. Statements (i) and (ii) are compatible as long as you yourself do not know that (ii) obtains.) On the other hand, the two-envelope paradox takes on even more bite if we build in (1) and (2), because of the seductiveness of the line of reasoning I am about to describe. But note well: statement (1) says only that the *actual* unknown amount in your envelope is such that, were you to know what it is, then you would still consider it equally likely that the other envelope contains twice or half that amount. This is vastly weaker than the implausible assumption that for *any* amount x, if x were in your envelope and you were to learn what amount it is, then you would still consider it equally likely that the envelope contains half or twice that amount.

[16]At any rate, you know this if we build into the scenario these two assumptions, in addition to (1) and (2) from the preceding note: (3) the amount y in the other envelope is such that if you were to learn what it is, then you still would consider it equally likely that the amount in yours is either half or twice y; and (4) you know this. And we really should build in these suppositions, if we build in the corresponding ones about x—because the decision scenario is supposed to be one in which the agent's information about the two envelopes is symmetrical in all relevant respects.

[17]To say that several standard probability assignments are "rationally eligible" does not mean, of course, that each of them is one that the agent is rationally permitted to adopt; rather, essentially it means that none of them conflict with the total available information. Insofar as they are all equally rationally eligible, it would be rationally inappropriate to adopt any one of them, over against the others.

[18]E.g., Cargile 1992: 212–13, Jackson *et al.* 1994: 44–45, and McGrew *et al.* 1997: 29.

[19]Letting $q_1,\ldots q_n$ be possible quantities in M,

$$U(\text{Stick}) = \Sigma\,[\text{pr}(\text{M contains } q_j)\cdot q_j]$$

Under the conditions of the problem, and given that the probability assignment satisfies the asymmetrical symmetry condition,

$$\begin{aligned} U(\text{Switch}) &= \Sigma\,\{[\text{pr}(\text{O contains } 1/2q_j \,\&\, \text{M contains } q_j)\cdot 1/2q_j] + [\text{pr}(\text{O contains } 2q_j \,\&\\ &\qquad \text{M contains } q_j)\cdot 2q_j]\} \\ &= \Sigma\,\{[1/2\text{pr}(\text{M contains } q_j)\cdot 1/2q_j] + [1/2\text{pr}(\text{M contains } q_j)\cdot 2q_j]\} \\ &= \Sigma\,[1.25\cdot\text{pr}(\text{M contains } q_j)\cdot q_j] \\ &= 1.25\cdot\Sigma\,[\text{pr}(\text{M contains } q_j)\cdot q_j]. \end{aligned}$$

Hence, U(Switch) = 1.25·U(Stick).

[20]Saying that rationality "requires the maximization" of U^δ means more than saying that rationality requires choosing an available act that happens to have a maximal U^δ-value. It also means that having a maximal U^δ-value is itself a *reason* why U^δ-maximization is rationally obligatory. The idea is that U^δ accurately reflects the comparative rational worth (given the agent's available information) of the available acts.

[21]This normative principle applies only to decision problems for which there is at least one rationally eligible standard probability assignment to epistemically possible states of nature. Let me acknowledge a complication that I do not address in the present paper. Given that there is some known minimal quantity of money that could possibly be in M or in O, I believe it can be shown that for the variant of the two-envelope decision problem that explicitly builds in conditions (1)–(4) described in notes 15 and 16 above, there is no rationally eligible standard probability assignment to epistemically possible states of nature. (This is because of features structurally similar to those at work in the "surprise examination" paradox.) Thus, the normative principle just stated does not apply to this version of the problem. Let me also acknowledge a second complication: according to the principle as stated, rationality requires U^δ-maximization even if the fact that U^δ ranks the available acts according to O is merely an accidental coincidence. What is really wanted, however, is some feature of U^δ *guaranteeing* that U^δ ranks the acts according to O. In Horgan forthcoming I discuss both complications (plus two others that I have lately discovered), and I propose a modified normative principle designed to accommodate them.

[22]Notice that U^y and U^z both are well-defined notions for the coin-flipping situation (and for the coin-flipping variant of the urn case), and that for both U^y and U^z the relevant probabilities remain at 1/2—i.e., pr(M contains 1/2y) = pr(M contains 2y) = 1/2, and pr(M contains z and O contains 2z) = pr(M contains 2z and O contains z) = 1/2. These epistemic probabilities, for pairs of y-based and pairs of z-based state-descriptions, rest on considerations analogous to those in section 2.5 above; in the coin-flipping scenario, the asymmetrical symmetry condition yields a pair of coarse-grained, equi-probable, epistemic possibilities in much the same way that the symmetry condition does in the original scenario.

[23]I thank Robert Barnard, David Chalmers, David Henderson, Nenad Miscevic, David Shoemaker, John Tienson, Ruth Weintraub, and Paul Weirich for helpful comments and discussion.

References

Arntzenius, F. and D. McCarthy. (1997) The two envelope paradox and infinite expectations. *Analysis* 57: 42–50.

Broome, J. (1995) The two-envelope paradox. *Analysis* 55: 6–11.

Cargile, J. (1992) On a problem about probability and decision. *Analysis* 54: 211–16.

Castell, P. and D. Batens. (1994) The two-envelope paradox: The infinite case. *Analysis* 54: 46–49.

Chalmers, D. (Unpublished) The two-envelope paradox: A complete analysis?

Follesdal, D. (1967) Knowledge, identity, and existence. *Theoria* 33: 1–27.

Horgan, T. (Forthcoming) The two-envelope paradox and the foundations of rational decision theory. In *Rationality and Irrationality: Proceedings of the 23rd International Wittgenstein Symposium*, eds. B. Smith and B. Bogaard. Vienna: hpt-öbv.

Jeffrey, R. (1983) *The Logic of Decision*, second edition. Chicago: University of Chicago Press.

Jackson, F., P. Menzies, and G. Oppy. (1994) The two envelope 'paradox'. *Analysis* 54: 43–45.

Kaplan, D. (1969) Quantifying in. In *Words and Objections: Essays on the Work of W.V. Quine*, eds. D. Davidson and J. Hintikka. New York: Humanities Press.

Kaplan, D. (1978) Dthat. In *Syntax and Semantics 9: Pragmatics*. ed. P. Cole. New York: Academic Press.

McGrew, T., D. Shier and H. Silverstein. (1997) The two-envelope paradox resolved. *Analysis* 57: 28–33.

Nalebuff, B. (1989) The other person's envelope is always greener. *Journal of Economic Perspectives* 3: 171–81.

Scott, A. and M. Scott. (1997) What's in the two envelope paradox? *Analysis* 57: 34–41.

Quine, W. V. (1969) Reply to Kaplan. In *Words and Objections: Essays on the Work of W. V. Quine*, eds. D. Davidson and J. Hintikka. New York: Humanities Press.