Logique & Analyse **225** (2014), 45–57

DECIDING AND TIME: REFUSING DEVILISH OFFERS

KARIM ZAHIDI*

ABSTRACT

In Gracely's decision theoretic puzzle, the Devil's Offer, a subject is offered to take part in a lottery in which she stands to win eternal bliss or eternal torment. Her chances of winning the lottery increase as time elapses. Expected utility maximization seems to lead to the paradoxical situation in which the subject indefinitely postpones the lottery, resulting in the least desirable outcome. However, as various authors have pointed out with respect to decision problems similar to the Devil's Offer, the reasoning that leads to the paradoxical conclusion is not valid. In particular, distinguishing between a synchronic and diachronic version of decision problems in which the number of choice options is infinite, allows for a fine-grained analysis of the requirements of rational decision making. The aim of this paper is twofold. I will show that Arntzenius et al.'s (2004) analysis of problems that are structurally similar to Gracely's decision problem can be adapted to show that the latter can be handled within Bayesian decision theory. Furthermore, I will discuss a variation on the diachronic version of the Devil's Offer, in which the subject is unaware of the fact that she is confronted with infinitely many choices. This modified version seems to be a genuine case in which the principle of utility maximization leads to a paradoxical result. However, I will show that this paradox is only apparent and that the problem can also be resolved within a broadly Bayesian conception of rationality.

1. Introduction

Decision problems involving infinities are often perplexing: seemingly sound reasoning leads to irrational outcomes. One example of such a decision problem is the Devil's Offer (Gracely, 1988). In this decision problem, a subject can enter a lottery in which she can win eternal bliss or eternal torment and in which her chances of winning the lottery increase as time elapses. She has to decide on which day she will enter the lottery. The principle of utility maximization seems to lead to a paradox: since on any

* The author would like to thank Erik Myin and an anonymous referee for valuable comments that resulted in improvements in style as well as substance of the paper. The Research Foundation Flanders (FWO), through the project "Senses as tools. A philosophy of sense modalities" (project G.0.321.09.N.10), supported the author.

doi: 10.2143/LEA.225.0.3011355.

given day postponing the lottery dominates not postponing, utility maximization seems to lead to a situation in which the subject indefinitely postpones the lottery, thus remaining in hell forever. The Devil's Offer belongs to a family of problems in which the infinitude of choice options wreaks havoc on patterns of reasoning which are perfectly legitimate in finite scenarios. Other such problems include Arntzenius and McCarthy's (1997) paradox of Heaven and Hell, Arntzenius et al.'s (2004) "Satan's apple"-problem, Pollock's (1983) "Ever Tasting Better Wine"-problem, and Landesman's (1995) "Terminating Trust"-problem. Problems such as these challenge our intuitions about rationality and question the idea that in these scenarios rational responses are available.[1] Although no solution to these problems is generally accepted, Arntzenius et al. (2004) show that some of these problems can be handled within a broadly Bayesian conception of rationality.[2] In particular they argue that since the expected utility does not attain a maximum, there is no point in equating rationality with maximizing expected utility.

Following Arntzenius et al. (2004)'s analysis of Satan's Apple-problem, I will show that by distinguishing between a synchronic and a diachronic version of Gracely's decision puzzle the problem can be handled within the Bayesian conception of rationality without leading to absurd conclusions (Section 2). Briefly, the distinction between diachronic and synchronic versions turns on the temporal order in which the subject makes the relevant decisions. In the synchronic version the subject decides on the first day which of the future offers she is going to accept and to reject. In the diachronic case she makes the decisions on a day to day basis, i.e. on the first day she decides whether to accept the Devil's first offer, leaving it open whether she's going to accept or reject the Devil's offer on subsequent days. That distinction turns out to be crucial for the resolution of the problem. In both cases it is important that the subject knows that the number of choices she has to make is potentially infinite and that expected utility does not attain a maximum over the available choice profiles. (A choice profile is the complete sequence of choices the subject makes. Since the number of choices she will be making is potentially infinite, the number of choice profiles is infinite.)

[1] Bermudéz (2009) distinguishes three different explanatory projects in which the concept (and hence a theory of) rationality plays a role: (a) the project of action-guiding, (b) the project of normatively assessing actions and (c) the project of explaining and predicting actions. In what follows, I will only focus on the first dimension, i.e. I will only be interested in how far rationality is suitable for guiding actions.

[2] The Bayesian conception of rationality assumes that when faced with a decision problem, a rational subject is able to rank the possible outcomes of her actions according to utility. This ranking should obey a number of criteria (e.g. the ranking should be transitive). Given such a ranking a rational decision is one which results in the outcome with the maximal utility.

In Section 3 I introduce a modified version of the diachronic Devil's Offer, in which the subject is unaware of the fact that the number of available choice profiles is infinite.[3] More precisely, on each day the Devil offers the subject a choice between holding the lottery on that day or postponing the lottery till the next day, thus improving the odds of winning the lottery. On each day the subject is presented only with two choice options and, in contrast with the original diachronic version, she is unaware of the fact that if she decides to postpone the lottery, the Devil will make her a similar offer on the subsequent day. Thus, the number of available choice profiles is infinite. In contrast with the original version, the subject is unaware of the latter fact and hence she has reasons to believe that expected utility can be maximized in this case. But, following this strategy requires her to postpone the lottery on each day, leading to the least desirable outcome (with probability 1) of remaining in hell for eternity.

However, in 3.2 I will show that the modified version can be handled within a decision theoretic framework, by taking into account the fact that evaluating expected utility has to be done under uncertainty. To incorporate that uncertainty in the subject's model one can consider the problem as a two-person's game in which at each stage, the Devil offers to postpone the lottery and the subject has to decide whether to accept the Devil's offer or not. But, in contrast with the original version, the subject is unaware that this game has (potentially) infinitely many rounds. Whereas at the start of the game it is rational to accept the Devil's offer, as the game unfolds and as the Devil's strategy becomes clear, the subject should consider whether her initial strategy of accepting the offers made by the Devil can maximize her chances of escaping from hell. To do this, she can compare the expected utility of her initial strategy with the expected utility of abandoning that strategy. A close analysis shows that after a number of rounds in the game, it is rational for the subject to assume that the number of available choice profiles is infinite and that her initial strategy will lead to the least desirable outcome of remaining in hell for eternity. Once the subject is aware of that fact, the problems that she faces in choosing the right course of action are similar to the problems in the original version.

2. The Devil's Offer[4]

Mrs C dies and goes to hell. She is approached by the Devil with the following offer. She is allowed to take part in a lottery in which the prize is

[3] In discussing Gracely's original problem, Priest (2000) actually gives a version of the problem which closely resembles the modified version.

[4] This section is based on the analysis of the Satan's Apple Problem by Arntzenius et al. (2004).

spending eternity in heaven. If she loses, she remains in hell forever. The lottery will be organised only once, but Mrs C is allowed to select the day on which it is to take place. If she decides that the draw is to take place on the r-th day, her chances of winning will be $1 - \frac{1}{r+1}$. When should Mrs C take her chance? Suppose that Mrs C adopts a utility-based approach and reasons as follows: by spending an extra day in hell, I will suffer a finite amount of torment (finite negative utility). However, this is largely compensated by the fact that spending another day in hell strictly monotonically raises my chances of gaining infinite bliss — i.e. the expected utility of having the draw tomorrow is higher than the expected utility of having the draw today. It therefore surely makes sense to wait another day. But this kind of reasoning is valid on every day; so it seems that trying to maximize expected utility, leads to the paradoxical situation in which the draw is indefinitely postponed and Mrs C stays in hell forever.

As Arntzenius et al. (2004) point out, decision problems in which the number of choice profiles is infinite come in two versions: a diachronic and a synchronic version. In the synchronic version Mrs C decides once and for all which of the Devil's offers she is going to accept and to reject, i.e. on the first day in hell, she decides which of the future offers she is going to accept and to reject.[5] In the diachronic case she makes the decisions on a day to day basis, i.e. on the first day she decides whether to accept the Devil's first offer, leaving it open whether she's going to accept or reject the Devil's offer on subsequent days. In other words, in the synchronic version Mrs C decides on her complete choice profile on the first day, whereas in the diachronic version the choice profile is built up on a day to day basis.

In the synchronic case, the reasoning that leads to accepting all the offers made by the Devil is not valid. To see this, note that when Mrs C decides whether to postpone the lottery on the day r, she reasons that postponing on day r is the dominant choice, regardless of the choices she makes about the subsequent offers. Hence, she concludes that her optimal choice profile, has to involve accepting the Devil's offer on day r. While this is true in case the sequence of choices is finite — i.e. if the number of available choice profiles is finite — it need not be true if the sequence is potentially infinite. For in the latter case it might very well be that the expected utility function does not attain a maximum. Hence, while postponing the lottery on day r is the dominant option, it does not follow that including that option in one's choice profile will result in a sequence of choices that maximizes expected utility.

[5] Mrs C *accepts* the Devil's offer on a given day when she decides to *postpone* the lottery; Mrs C *rejects* the Devil's offer when she decides to hold the lottery on that day.

Mrs C is thus not rationally required to accept all of the Devil's offers, but the question remains which offers she should accept. There is no univocal rational answer to the question on what day the lottery should be held. For any choice, there are infinitely many choices that would result in a larger expected utility. However, Mrs C is required to refuse one offer, for not refusing one offer will leave her in hell for ever. The only guidance, based on expected utility calculations, is that it is rational to avoid an outcome that she finds least desirable, i.e. the outcome for which expected utility is not higher than a treshold she has fixed. Note however that she can approximate the upper bound on the expected utility $EU(C_r)$ to any degree of accuracy she chooses.

In the diachronic version what the rational course of action is, depends on two factors: "it depends on whether she can bind herself to future courses of action" and on the manner in which she expects that her decision about the Devil's offer on a given day will influence her decisions about subsequent offers (Arntzenius et al., 2004, p. 265). If Mrs C is able to bind herself to a course of action on which she decides on the first day, then her future decisions are completely determined. In this case the diachronic case reduces to the synchronic case: Mrs C decides on the first day what her complete profile will be. Of course, some agents lack the capacity for binding. In particular it might be the case that Mrs C's decision of to accept the Devil's offer on day *r* influences her decisions on subsequent days. Suppose for example that Mrs C knows that if she decides to postpone on the first day, she will postpone on all other days, then it is rational for her to hold the lottery on the first day. Of course, Mrs C might lack the capacity to bind and believe that her choices are independent.[6] In that case, rationality requires her to accept all of the Devil's offers. This might seem an unwelcome conclusion. However, Arntzenius et al. (2004) argue convincingly that this is indeed required by rationality.[7]

[6] It might be objected that Mrs C's choices are not independent: deciding to reject the Devil's offer on day *r* implies that the game ends on day *r* and no further choices need to be made. However, the choices can be independent in the sense that if Mrs C's choice on a given day is such that there are choices left to be made, then these choices are independent of that choice.

[7] Briefly, their argument runs as follows. If Mrs C really believes that her decisions are independent, then it is as if for each decision she consults a separate expert. The expert who decides on whether to accept the Devil's offer on the first day is certain that his choice will not influence the other experts. This independence makes it impossible for the experts to rationally coordinate their actions. Hence the only option they have, is to accept the offer.

3. The Devil's offer under uncertainty

3.1. *Formulation of the problem*

I now turn to a modified version of the Gracely's problem. Mrs C's good friend, Mrs D, dies and finds herself in hell. On her first day in hell she is approached by the Devil with the following offer. She is allowed to take part in a lottery in which the prize is spending eternity in heaven. If she loses the lottery, she remains in hell for ever. She can decide to hold the lottery on her first day or she can decide to postpone the lottery until the next day, thus increasing her chances of winning from $\frac{1}{2}$ to $\frac{2}{3}$. Since the expected utility of holding the lottery on the second day is higher than the expected utility of holding the lottery on the first day, Mrs D decides to postpone the lottery till the second day. On her second day in hell, the Devil approaches Mrs D and tells her that they can have the lottery on the second day or, if Mrs D is willing to spend another day in hell, they can have the lottery on the third day, thus increasing the chances of winning from $\frac{2}{3}$ to $\frac{3}{4}$. As long as Mrs D accepts the Devil's offer, the Devil continues to offer her a similar deal on the subsequent day: holding the lottery on day r with winning chances $1 - \frac{1}{r+1}$ or holding the lottery on day $r + 1$ with winning chances $1 - \frac{1}{r+2}$.

The crucial difference between the original diachronic Devil's Offer and the modified version is that, while Mrs C knows that she is facing a decision problem in which the number of choice profiles is infinite, Mrs D is ignorant about that latter fact. This difference implies that the resolution to the original problem as discussed in Section 2, is not applicable to the modified version. It fails because in the original problem Mrs C is aware of the fact that she is facing a decision problem in which the number of choice profiles is infinite, and that her expected utility function does not attain a maximum over this infinite set. This implies that maximizing expected utility is impossible. Hence, rationality does not require her to choose the dominant option at each stage of the game. In the modified version, however, Mrs D is unaware of the fact that the number of choice profiles is infinite and therefore she may think that the expected utility function can in fact be maximized over the relevant choice profiles. It thus seems that in the modified version of the problem, Mrs D's rational choice option is to accept the Devil's offer whenever it is made, thus indefinitely postponing the lottery. Is the modified version a genuine example of a problem where rational analysis breaks down? I will argue that the answer to this question is negative and will show that Mrs D can analyze the problem in such a way that rationality does not require that she postpones the lottery indefinitely. The key point is that, as the game unfolds, Mrs D can, based on her experience in the previous stages of the game, come to see that her initial assumption (that there are only finitely many available choice profiles) might be wrong. Once she realizes this fact, she has to take into account that there might actually be infinitely many choice profiles from which she has to choose, and hence that she is not required to choose the dominant option at each stage in the game.

3.2. *Resolving the problem*

In each round Mrs D is offered two alternatives: postponing or not postponing the lottery. In the first few rounds this choice seems to be equivalent to the choice between having the lottery today or tomorrow: she is unaware of the fact that in subsequent rounds the Devil will continue to offer her better and better odds in the lottery. Consequently, on each day Mrs D compares the expected utility of holding the lottery today with the expected utility of holding the lottery on the next day. In other words, on day r, Mrs D believes that her choice is restricted to two different options: having the lottery today or having the lottery tomorrow. Since Mrs D thinks that there are only finitely many choice profiles to choose from, she chooses the dominant option, i.e. she accepts the Devil's offer. After a few rounds, as she starts to understand that the Devil might continue to offer increased chances of winning in the lottery when postponing, she realizes that she might be confronted with the choice between infinitely many choice profiles, none of which maximizes expected utility. She should thus take into account utility maximization might be impossible. However, from the information Mrs D has access to, it is not clear whether and when she should stop trying to maximize expected utility.

In order to simplify the mathematical modeling, I will assume that spending a finite amount in heaven or hell has no utility; that the utility of spending eternity in heaven — denoted by U_H — is finite and positive and, that the utility of spending eternity in hell — denoted by U_h — is finite and negative. These assumptions do not alter the paradoxical nature of the problem, since, as we have seen in the discussion in Section 2, the latter is solely due to the fact that the expected utility function does not attain a maximum over the infinite set of available choice profiles.

Returning to Mrs D's problem as to whether and when she should give up her initial strategy, one way to decide that issue is to calculate the expected utility of continuing with her strategy. Let $EU(S_r)$ denote the expected utility, on day r, of continuing with the initial strategy. Evaluating this expected utility has to be done under incomplete information: Mrs D can only guess what the Devil's move will be in the next rounds. I.e. Mrs D's calculation of the expected utility of not abandoning the initial strategy on the day r, must involve a subjective assessment of the likelihood that the Devil is playing out a strategy that might result in her having a choice between infinitely many choice profiles, in which case her strategy of trying to maximize expected utility leads to indefinitely postponing the lottery.

Mathematically this is expressed by the fact that $EU(S_r)$ is the weighted sum of the expected utility of having the draw on some future day and the utility of not having the draw at all, i.e. staying in hell for eternity. This weighted sum is of the following form:[8]

$$EU(S_r) = \sum_{k>r} q_{(r,k)} EU(C_k) + p_r U_h \tag{1}$$

Here $EU(C_k)$ denotes the expected utility of the choice profile

$$C_k = (\underbrace{\text{accept, accept, …, accept}}_{k-1}, \text{ reject}),$$

i.e. of holding the lottery on the k-th day. It is easy to see that $EU(C_k)$ is given by the following expression:

$$EU(C_k) = \left(1 - \frac{1}{k+1}\right) U_H + \frac{1}{k+1} U_h. \tag{2}$$

The coefficients $q_{(r,k)}$, $k = r+1, r+2, \ldots$ are numerical representations of Mrs D's degree of belief, on day r, in the proposition that the lottery will take place on the k-th day given the fact she continues with her initial strategy.[9] The coefficient p_r represents Mrs D's degree of belief, on day r, in the proposition that the draw will never take place, under the condition that she sticks to the initial strategy. Adopting a broadly Bayesian approach, these degrees of belief are in fact subjective probabilities attached to mutually exclusive and exhaustive events and thus sum to one, i.e. for each value of r:

$$p_r + \sum_{k>r} q_{(r,k)} = 1. \tag{3}$$

Furthermore, since at stage r in the game, Mrs D only knows that the Devil has offered her to defer the lottery one more day, it is reasonable for Mrs D

[8] One may wonder whether the term "expected utility" as used in this paragraph is apt. Certainly the term is used in a non-standard way. Expected utilities usually depend on objective probabilities whereas the quantity $EU(S_r)$ depends on subjective probabilities. For want of a better term I will refer to this quantity as "expected utility". Note also that formally this weighted sum looks like the expected utility of a mixed strategy (Hájek, 2003). However it is important to note that it represents the expected utility of a pure strategy.

[9] Note that (1) is an infinite sum and thus there is no guarantee that the sum converges to a finite value. However this is not problematic for what follows. What is important to note is the fact that if it converges, it converges absolutely. Hence any rearrangement of the terms will result in the same finite or infinite sum.

to attach greater credence to the proposition that the lottery will take place at stage $r+1$ than to the propositions expressing the fact that it will take place on some future day, i.e.

$$q_{(r,r+1)} > q_{(r,r+2)} > q_{(r,r+3)} > \dots .$$

The subjective probabilities $q_{(r,k)}$ and p_r reflect Mrs D's assessment, on the r-th day, of what the Devil might do in the following stages of the game and thus depend on the behaviour of the Devil in the previous rounds. Therefore, they vary from one round to the next, i.e. they are functions of r. At the start of the game, i.e. for small values of r, Mrs D will reasonably assume that p_r is close to zero, and that the probabilities $q_{(r,k)}$ are also very small for values of k larger than $r+1$.[10] This reflects her — at that moment justified — belief that it is highly unlikely that the Devil will repeat his offer beyond stage $r+1$.[11] As the game unfolds and she gathers more information about the Devil's strategy, based on his behaviour in the previous stages of the game, she begins to doubt whether her assessment of the Devil's strategy is correct. This will result in a revision of the relevant subjective probabilities. I will assume that it is rational for Mrs D to update the subjective probabilities in such a way that the sequence $p_1, p_2, p_3 \dots$ forms an increasing sequence approaching one, i.e.

$$\lim_{r \to \infty} p_r = 1. \tag{4}$$

Given this assumption, which will be defended in Section 3.3, it follows from (3) that for large values of r, p_r will be much larger than the probabilities $q_{(r,k)}$, reflecting her doubts that the initial strategy will ever result in holding the lottery. This implies that for sufficiently large values of r, it is no longer the case that in the expression (1) for $EU(S_r)$ the term involving $EU(C_{r+1})$ (or any other term involving $EU(C_{r+2})$, $EU(C_{r+3})$, …) outweighs all other terms, i.e. the term $p_r U_h$ becomes dominant for large values of r. Since U_h is negative, it is no longer clear whether for sufficiently large values of r, the expected utility of continuing with her initial strategy is non-negative.

[10] I assume that Mrs D is a perfect Bayesian rational agent implying that she considers all possible scenario's. This may seem unrealistic: perhaps at the start of the game she does not think of the scenario where the Devil keeps proposing to postpone the lottery as a possibility. This implies that she does not assign any probability to that scenario. Since my aim is to show that the problem can be solved within the Bayesian conception of rationality, I will henceforth assume that Mrs D is a perfect Bayesian rational agent (see footnote 12 for some further comments on the non-perfect case).

[11] For small values of r, $EU(S_r)$ is almost equal to $EU(C_{r+1})$.

To elaborate this last point, I will assume that the utilities of spending eternity in hell and heaven cancel out, i.e. $U_h = -U_H$ (with $U_H > 0$). The expression for $EU(C_k)$, i.e. expression (2), can then be rewritten as $EU(C_k) = (1 - \frac{2}{k+1})U_H$ (note that for any positive value of k, $EU(C_k)$ is positive). Substituting this expression in (1) and simple algebraic manipulation yields

$$EU(S_r) = \left(\sum_{k>r} q_{(r,k)} - \sum_{k>r} q_{(r,k)} \frac{2}{k+1} - p_r\right) U_H. \tag{5}$$

Since (3) implies $\sum_{k>r} q_{(r,k)} = 1 - p_r$, (5) is equivalent with

$$EU(S_r) = \left(1 - 2p_r - \sum_{k>r} q_{(r,k)} \frac{2}{k+1}\right) U_H. \tag{6}$$

It follows immediately from (6) that $EU(S_r)$ is negative when p_r is larger than the threshold value 1/2. Given the assumption that the sequence p_1, p_2, p_3, … approaches one, there exists a value for r such that p_r exceeds the threshold. For this value of r, the expected utility of continuing with her initial strategy on the r-th day is negative. Given the fact that there are other choices which have positive expected utility, e.g. having the lottery on the r-th day, she should abandon her initial strategy. It is possible that already on an earlier date the choice of holding the lottery on that day dominates her initial strategy. Whether this will be the case depends on the exact numerical values of the parameters p_r and $q_{(r,k)}$. Although this shows that by rationality, she is not obliged to continue with her initial strategy, this does not answer the question when she should participate in the lottery.

Before I address this question, it is useful to point out another interpretation of the probabilities p_r and $q_{(r,k)}$. Recall that p_r is a measure for the degree of belief that Mrs D attaches to the proposition that her initial strategy will not maximize her chances to get her out of hell. Hence, as p_r increases, Mrs D's confidence in her strategy becomes smaller. Since Mrs D's initial strategy consists of choosing the dominant option — i.e. postponing the lottery — applying the principle of utility at each decision step. Hence, p_r is a measure of Mrs D's confidence in the applicability of the principle of utility maximization when dealing with the Devil's offer. But, as already pointed out, the inapplicability of that principle is connected with the fact that the number of choice profiles is infinite. Hence, the probability p_r can be interpreted as a measure of Mrs D's degree of belief (at stage r of the game) in the fact that there are infinitely many choice profiles. Similarly, the probability $q_{(r,k)}$ can be interpreted as a measure of Mrs D's degree of belief (at stage r of the game) in the fact that the number of available choice profiles is k.

Given this reinterpretation of the probability p_r, the significance of the threshold 1/2 becomes clear. If at a given stage in the game, Mrs D thinks it is more likely that she is in fact confronted with infinitely many choice profiles rather than with finitely many — i.e. if p_r is larger than the threshold — she may as well assume that her predicament is akin to the predicament of her good friend Mrs C in the original (diachronic) version of the problem. This implies that the analysis given in the previous section becomes applicable. In particular it follows that Mrs D is not required to accept all the Devil's offers, unless she lacks the capacity to bind and furthermore believes that her decision to accept or reject the Devil's offer on a given day does not influence her subsequent decisions. In the latter case, she is rationally bound to accept all of the Devil's offers, thus remaining in hell forever.

3.3. *Defending condition (4)*

The resolution of the paradox in 3.2 turns crucially on the assumption that the sequence of probabilities p_1, p_2, p_3, … approaches one. Absent this assumption, there is no guarantee that there exists a value of r such that p_r exceeds the threshold value. Consequently, there is no guarantee that for some finite r-value Mrs D should abandon her initial strategy. Is this assumption justified? Recall that p_r is a measure for Mrs D's confidence in her initial strategy: a high confidence in the strategy will result in low values of p_r, while a lack of confidence is reflected by high values of p_r. At the start of the game, Mrs D is justified in having high confidence in her initial strategy. After a few rounds however, this initially justified confidence becomes problematic: as she comes to understand the Devil's strategy, it becomes less and less clear that her initial belief in the strategy is still justified. She should therefore adjust the p_r-values in accordance with her loss in confidence, i.e. p_{r+1} should be larger than p_r.

This is borne out by the fact that, ex hypothesi, the degrees of belief are probabilities and hence the updates should accord with Bayes' theorem, which in this case yields the following:

$$p_{r+1} = \frac{p_r}{1 - q_{(r-1,\, r)}} \tag{7}$$

Since on day $r-1$ the Devil did renew his offer, either Mrs D did or did not believe that there was a chance that the lottery could take place on day r. In the former case $q_{(r-1,\, r)} > 0$, which yields $p_{r+1} > p_r$. In the latter case $q_{(r-1,\, r)} = 0$, but then also $q_{(r-1,\, k)} = 0$ for values of k larger than r, hence $p_{r-1} = 1$. Consequently Mrs D would already have held the lottery on day

$r-1$. Hence, the sequence of p_r-probabilities forms a strictly increasing sequence when the probabilities are updated according to Bayes' theorem. Furthermore, numerical simulation shows that the sequence of p_r-probabilities is unbounded within the unit interval, i.e. can dominate any value smaller than 1.

An alternative argument for condition (4) can also be made relying on the alternative interpretation of the p_r-probability. Recall that the p_r-probabilities can also be interpreted as a measure for Mrs D's degree of belief in the fact that there are infinitely many choice options available. Now, as Mrs D advances in the game, and gathers more information about the Devil's strategy, she will notice that the number of choice profiles is larger than originally thought. On the first day, after the Devil makes his first offer, she comes to see that there are two choice profiles to choose from. On the second day, after the Devil made the second offer however, she can, in retrospect, see that she was mistaken about her belief that on day one there were only two choice profiles; in fact on the second day she realizes that there were at least three choice profiles. As the game unfolds, at each stage in the game, she retrospectively comes to see that in previous rounds she misjudged the number of available choice profiles. As the number of rounds increases, she is confronted with the fact that any reasonable guess about the number of choice profiles proved to be mistaken. Given this fact, it is rational to adjust the p_r-probabilities.[12]

4. Conclusion

The modified version of the Devil's Offer at first seems to be a puzzle that cannot be resolved within the Bayesian conception of rationality. The subject, who is confronted with the modified problem, lacks knowledge about the game, which is present in the original scenario. Since she is unaware of the fact that the number of choice profiles is infinite, it is rational for her to assume that she is required to choose the dominant option at each stage of the game. However, I have shown that from a certain stage on, which

[12] If Mrs D is not a perfect Bayesian rational agent, she might not assign any value to p_r (see footnote 10). This implies that there is no starting point to get the Bayesian belief revision process off the ground. However, the argument I have developed in this paragraph to defend condition (4), can be used by an agent who's not a perfect Bayesian rational agent, to come to see that it is rational to expand the space of possibilities. Indeed, while for a perfect Bayesian rational agent the space of possibilities is fixed throughout the game, this is not necessarily the case for a non-perfect agent. However, the argument given here convincingly shows that adapting the space of possible scenario's — i.e. scenario's that one takes into account in the deliberation of what to do — is the rational thing to do. Once the scenario in which the Devil keeps repeating his offer is considered as a possibility Mrs D can assign a probability to it and proceed with the Bayesian analysis.

can be determined by expected utility calculations, it is rational for the subject to assume that she is in fact confronted with infinitely many choice profiles. Once this stage is attained, the problem reduces to the original version of the problem. In particular this implies that what the rational course of action is, depends on a number of other factors: capacity for binding and the relationship between subsequent choices. If the subject has the capacity to bind, she can decide on a certain choice profile and stick to this profile. In choosing the choice profile, she cannot rely on the principle of utility maximization, since her utility function does not attain a maximum over the available choice profiles. She can only indicate which lower bound on utility she is willing to accept. If she lacks the capacity to bind, then what she should do, depends on how she thinks rejecting or accepting the Devil's offer on a certain day will influence her subsequent choices.

References

F. Arntzenius and D. McCarthy (1997). The two envelope paradox and infinite expectations. *Analysis*, 57:28-34.

F. Arntzenius, A. Elga, and J. Hawthorne (2004). Bayesianism, Infinite Decisions, and Binding. *Mind*, 113:251-283.

J.L. Bermudéz (2009). Decision Theory and Rationality. Oxford University Press, Oxford.

E.J. Gracely (1988). Playing games with eternity: The Devil's offer. *Analysis*, 48:113.

A. Hájek (2003). Waging War on Pascal's Wager. *The Philosophical Review*, 112(1):27-56.

C. Landesman (1995). When to Terminate a Charitable Trust. *Analysis*, 55:12-13.

J.L. Pollock (1983). How do you Maximize Expectation Value. *Noûs*, 17:409-421.

G. Priest (2000). *Logic. A very short introduction*. Oxford University Press, Oxford.

Karim Zahidi
Department of Philosophy
University of Antwerp
Prinsstraat 13
B-2000 Antwerpen
Belgium
E-mail: karim.zahidi@uantwerpen.be