RJ10493 (ALM1111-005) November 7, 2011
Mathematics

IBM Research Report

Reexamination of Allais' Paradox, Prospect Theory, and Expected Utility Theory

Nimrod Megiddo
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099
USA

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Reexamination of Allais' Paradox, Prospect Theory, and Expected Utility Theory

Nimrod Megiddo*

October 28, 2011

Abstract

The axioms of von-Neumann and Morgenstern's utility theory are not contradicted by experiments if the domain of outcomes is constructed properly. Preferences over outcomes take into account issues of regret with respect to alternative outcomes, and surprise with respect to prior probabilities of outcomes. Utilities can be derived as fixed points of a relations between the intrinsic value of an outcome and the money-equivalent value of a probability distribution over outcomes.

1 Introduction

Allais' Paradox [1, 2] is an empirical finding, which is often used for questioning the validity of von Neumann - Morgenstern's utility theory [6], as well as for justifying the so-called prospect theory [2]. Kahneman and Tversky [2] describe empirical results of preferences reported by humans, which seem to violate the expected-utility property of von Neumann and Morgenstern's utility functions. However, as we demonstrate in this paper, because in every decision problem, an "outcome" entails more than the mere amounts of money awarded to the players, a careful application of utility theory requires more than just associating utility values with amounts of money. Our long-term goal is not necessarily to refute Kahneman and Tversky's assertions about human behavior. We rather seek to distinguish elements of human preferences that persist even after attempts to coach individuals to follow more sound methodologies, from ones that can be modified by careful analysis or introspection.

2 Outcomes of decision problems

Because preferences of human beings are complex and involve emotions and feelings, the notion of an "outcome" of a decision problem has to be analyzed with care. In particular, the satisfaction of an individual from a certain outcome may depend not only on the amount of money he receives but also on the prior probability of the outcome, the alternative outcomes, and their respective prior probabilities. The outcomes also include the rewards to the other players. Hence, the satisfaction of a player from an outcome depends also on the rewards to other players, not only in the considered outcome but also in alternative outcomes.

Consider, for example, the decision problem depicted in Figure 1. In this problem, the decision maker decides whether or not to quit with zero payoff, or to let a coin be tossed. However, the result of the coin toss not revealed to the decision maker until after he has made a second decision. After the coin has been tossed,

*IBM Almaden Research Center, San Jose, CA 95120. email: megiddo@us.ibm.com

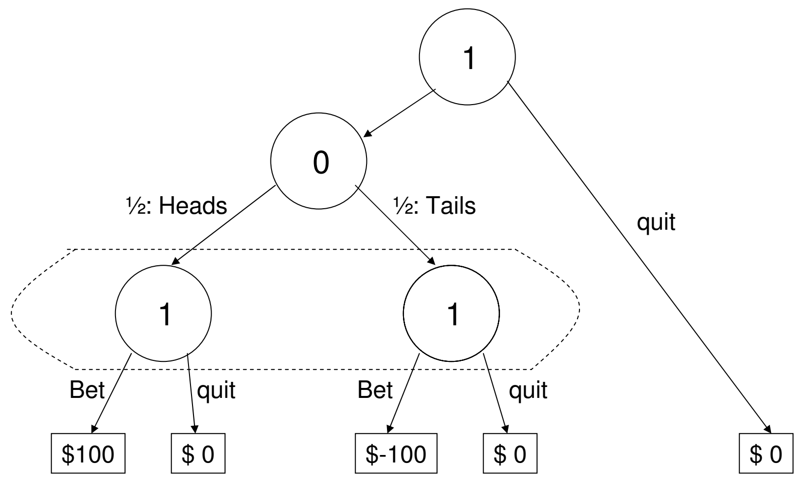

Figure 1: Multiple utility values for zero

the decision maker then has a second chance to quit with a zero payoff. If, again, he does not quit, then his gain is \$100 if the toss came out Heads, or \$-100 if the toss came out Tails. A naïve look at this problem suggests that there are only three outcomes, namely, \$0, \$100 and \$-100. From this naïve point of view, the question is essentially whether the decision maker prefers to quit rather gamble on \$100 versus \$-100. In particular, it implies that the decision maker is indifferent between quitting at the first chance and quitting at the second chance. A more careful look, however, suggests that the three outcomes with the same zero payoff are quite different from the point of view of emotions. Suppose the player first chooses to have the coin tossed, but then quits the game. If he later finds that the toss came out Heads, then he regrets that he did not gamble. On the other hand, if he later finds that toss came out Tails, then he is content that he did not gamble. If the player quits at the first chance, he never knows exactly what would have happened had he not quit. Thus, it seems that, at least for some players, the zero payoff after Tails is the most preferred, followed by zero payoff after a quit at the first chance, and least preferred zero payoff is the one after Heads.

The conclusion is that every leaf of a decision tree should be considered as a different outcome, so that the set of leaves, rather than the set of possible monetary payoffs, should be the domain of the utility function.

3 Allais' examples and prospect theory

Kahneman and Tversky [2] presented a pair of choice problems as a variation on Allais' example, where the extremely large gains in Allais' original example were replaced by moderate ones. Problems 1 and 2 are depicted in Figures 2 and 3 Kahneman and Tversky reported that 82% of the people preferred a sure pay of \$2400 over the gamble at the left branch in Problem 1, and 83% of the people preferred the bet on \$2500 with probability 0.33 to the bet on \$2400 with probability 0.34. They concluded that people's utility functions did not have the expected utility property, because there is no utility function $u : \{0, 2400, 2500\} :\to \Re$ consistent with such preferences that uses expected utilities to describe preferences over probability distributions.

One issue with the above argument against expected utility is that the two problems are presented in two very different situations. The standard utility theory has to be applied to the whole decision problem, where all the outcomes are enumerated. There can be several ways to resolve this issue, but they may have different consequences. For example, a random device may decide which problem to present to the player, or problem 2 may be presented after problem 1 (before or after the lottery has taken place), and so on. Figure

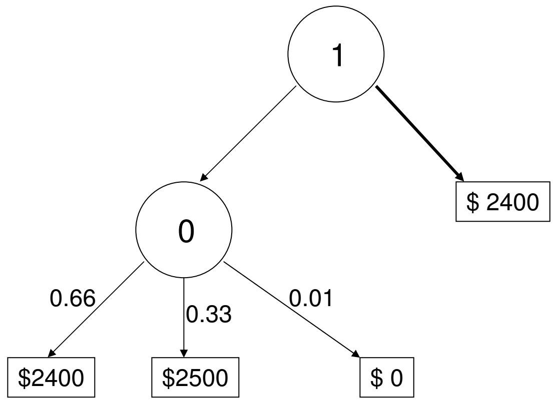

Figure 2: Problem 1

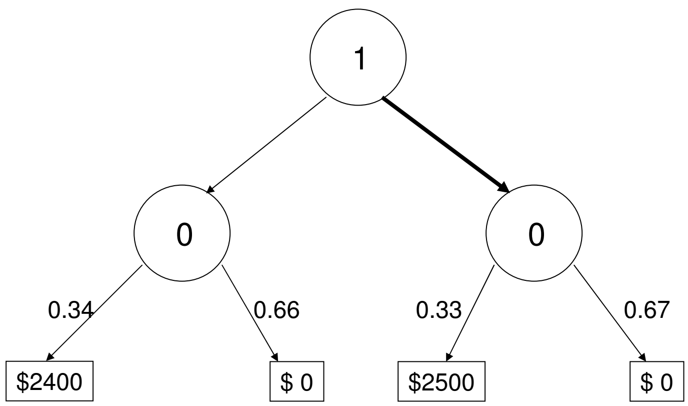

Figure 3: Problem 2

4 depicts a situation, where Player 2 decides which problem to present to Player 1, but the latter can form his preferences in advance.

Once the two problems are placed within the same decision tree, there is no justification to assume that the decision maker is indifferent between two outcomes with the same monetary payoff. This observation invalidates the proof of nonexistence of an expected-utility function. Moreover, the particular way in which the two problems are placed within one decision situation may affect the preferences. Suppose, for the moment, that the combined problem is the one depicted in Figure 4. Consider the outcome of zero payment that can happen with probability 0.01 if the Problem 1 is presented and the decision maker chooses to gamble. Compare it to the outcomes of zero payment that can happen with probabilities 0.67 and 0.66, respectively, if Problem 2 is presented. In problem 1, the player is very disappointed with the zero outcome because it was not anticipated to actually happen, whereas in Problem 2 it is much more likely to happen, so the disappointment is lesser. Similarly, the sure outcome of $2400 in Problem 1 is not as pleasing as the outcome of $2400 in Problem 2, whose prior probability is only 0.34, so the player may have a higher utility for the same dollar payoff of $2400 in Problem 2.

Kahneman and Tversky [2] also presented another pair of choice problems as follows. Problems 3 and 4 are depicted in Figures 5 and 6. Like in Problems 1 and 2, there are only three different monetary rewards

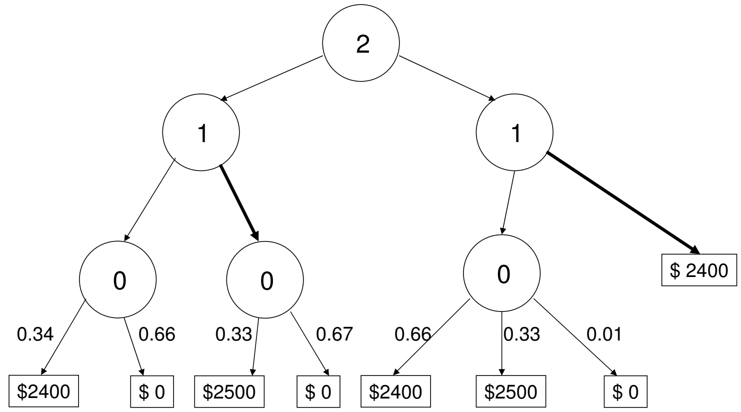

Figure 4: Another player decides which problem to present

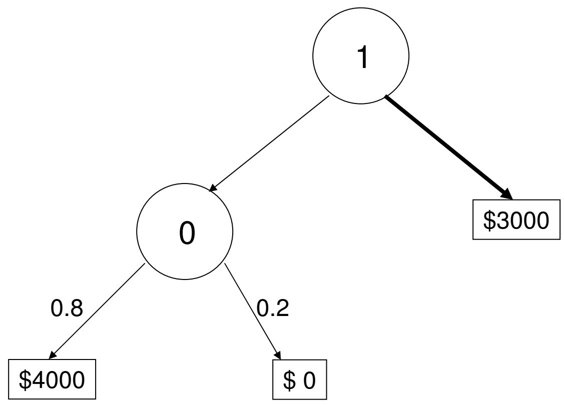

Figure 5: Problem 3

in Problems 3 and 4, namely, zero, \$3000 and \$4000. People who prefer in Problem 3 a sure payment of \$3000 and in Problem 4 the bet on \$4000 are claimed by Kahneman and Tversky not to have a vNM utility function. However, as we pointed out above, outcomes with the same monetary payoffs may have different utility values. For example, getting zero in problem 3 my be more disappointing than getting zero in Problem 4 because the prior probability of zero in Problem 3 is 0.2, and in Problem 4 it is 0.75 or 0.8, depending on the choice. Also, the payment of \$3000 in Problem 4 is more pleasing than the sure payment of \$3000 in Problem 3, and the payment of \$4000 in Problem 4 is more pleasing than the payment of \$4000 in problem 3.

The setting in which Kahneman and Tversky's theory was developed is that a subject is asked to choose between two alternatives, and then asked again to choose between two other alternatives. We do not address here the effect of the order in which these two pairs of alternatives are presented. To formulate the preferences of a decision maker, we simply consider a setting in which a disinterested Player 2 picks a probability distribution over rewards, and then Chance picks a reward for Player 1 according to that chosen distribution. Thus, such a setting with only possible distributions is depicted in Figure 7, where the payoffs are to Player 1. There are four possible outcomes in this process. Denote the set of these outcomes by D. If the preferences over probability distributions over D satisfy the axioms of von Neumann and Morgenstern,

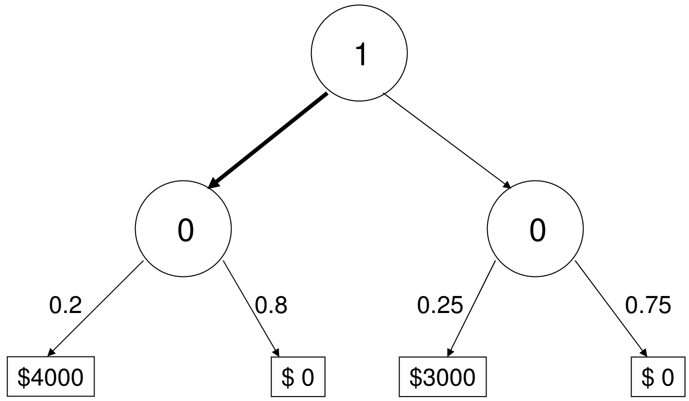

Figure 6: Problem 4

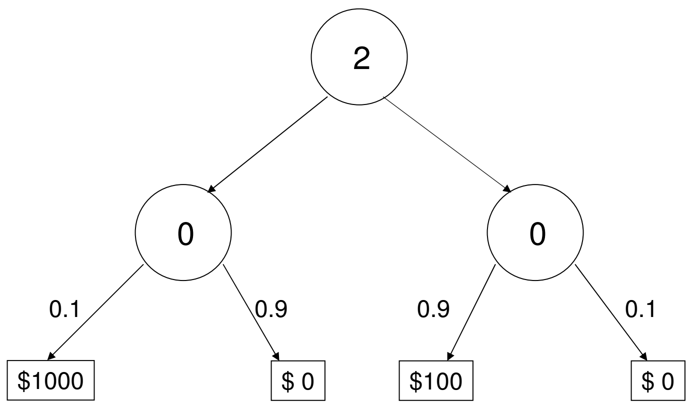

Figure 7: A simple setting

then they can be represented by a utility function $u : D \to [0, 1]$ that has the expected utility property. Note that the player may strictly prefer the outcome of zero payment in the left subtree over the zero payment in the right subtree. However, these two outcomes have the same value under the Kahneman - Tversky theory.

It is possible to assume that the four utilities $u(\$1000, 0.1)$, $u(\$100, 0.9)$, $u(\$0, 0.9)$ and $u(\$0, 0.1)$ are derived from a "pure" utility from money $u_0(m)$ and a weighting function $w(p)$, so that $u(m, p) = u_0(m)/w(p)$. Under this approach, if Player 2 picks the left subtree with probability p and the right subtree with probability $1-p$, then the final distribution of rewards is \$1000 with probability $0.1\,p$, zero with probability $0.9\,p$, another zero with probability $0.1(1-p)$ and \$100 with probability $0.1(1-p)$. Thus, the expected utility of Player 1 in this case is equal to

$$U = 0.1\,p\,u(1000)/w(0.1) + 0.9\,p\,u(0)/w(0.9) + 0.9(1-p)\,u(100)/w(0.9) + 0.1(1-p)\,u(0)/w(0.1)\ .$$

However, once the probability p has been fixed, the utilities of the outcomes become $u(1000)/w(0.1\,p)$, $u(0)/w(0.9\,p)$, $u(100)/w(0.9(1-p))$ and ,$u(0)/w(0.1(1-p))$, respectively, so we must have

$$U = 0.1\,p\,u(1000)/w(0.1\,p)+0.9\,p\,u(0)/w(0.9\,p)+0.9(1-p)\,u(100)/w(0.9(1-p))+0.1(1-p)\,u(0)/w(0.1(1-p))\ .$$

We could consider a possibility between these two theories, namely, preferences over probability distributions over the distinct pairs (m_i, p_i) that occur in the decision tree for any strategy of the decision maker. Kahneman and Tversky propose a utility function based on values $V(m_i)$ and a weighting function $\pi : [0, 1] \rightarrow \Re$ so that the utility from a lottery that yields m_i with probability p_i is equal to $\sum_i \pi(p_i) \cdot V(m_i)$. Now, if we adopt that view that the "value" of a monetary reward m_i whose prior probability is p_i is of the form $V(m_i)/\sigma(p_i)$, where $\sigma(\cdot)$ is another weighting function, then it is possible that the expected value $\sum_i p_i V(m_i)/\sigma(p_i)$ is the same as the Kahneman and Tversky value. This happen if $p\sigma(p) = \pi(p)$, i.e., $\sigma(p) = p/\pi(p)$. However, as we demonstrate later, there could be two outcomes i, j with $(m_i, p_i) = (m_j, p_j)$ but with different utilities and therefore different contributions to the utility from the distribution, so terms of the form $\pi(p_i) \cdot V(m_i)$ are insufficient.

4 Accounting for prior probabilities

The discussion above suggests that human decision makers may enjoy realized gains, when the gains have small prior probabilities, more than the same gains when they have larger prior probabilities. Similarly, humans may be disappointed with realized losses, when the losses have small prior probabilities, more than from the same losses when they have larger prior probabilities. In a decision tree with monetary payoffs, for a fixed strategy of the decision maker, distinct outcomes are represented by distinct leaves of the tree. For a fixed strategy, for each leaf i there is a monetary reward m_i and a probability p_i that the leaf will realize as the outcome. The von Neumann–Morgenstern theory applies to preferences over probability distributions over all possible outcomes. The Kahneman–Tversky theory applies to preferences over probability distributions over the set of distinct rewards. We now compare the two approaches through some examples, starting from very simple ones.

Example 1. Consider the very simple situation depicted in Figure 8. The choice is made by a random

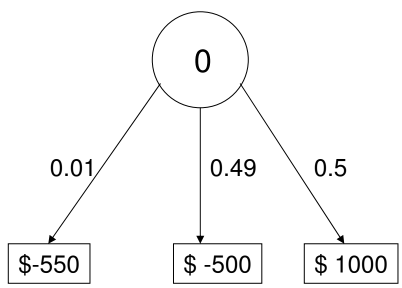

Figure 8: A simple lottery with no decisions

device, and the payoff is made to a passive player. The player may view this situation as very close to a 50:50 bet on gaining \$1000 or losing \$500. Thus, if the outcome turns out to be the less anticipated loss of \$550, the utility of losing \$550 in this circumstance may be lower than the utility of such a loss if it were more anticipated, for example, as in the situation depicted in Figure 9. However, with a normalized utility function $u(\cdot)$, we will have in the situation of Figure 8, $0 = u(\$-550) < u(\$-500) < u(\$1000) = 1$, and in the situation of Figure 9, $0 = u(\$-550) < u(\$-450) < u(\$1000) = 1$, but the value of $u(\$-450)$ could be *smaller* than that of $u(\$-500)$ due to the normalization. If the simple tree, as in the two situations above, reflects the entire set of outcomes, then there is no need to consider preferences over all possible probability distributions of outcomes other than the one built into the tree.

The above examples are quite simple because they require no decisions, so the preferences of the player do not have operational consequences. In particular, the issue of regret over one's own choice is not present.

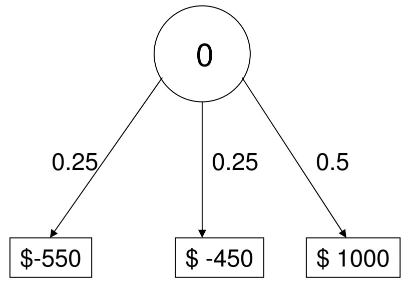

Figure 9: Another simple lottery with no decisions

However, if we ask the player to choose one of two distributions, then we obtain a more complicated situation as we explain below.

Example 2. Consider the situation depicted in Figure 10. To determine which subtree is preferred in

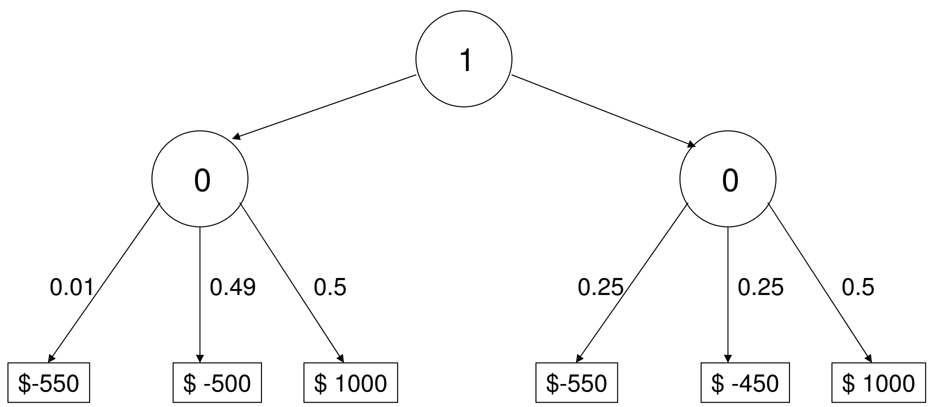

Figure 10: A simple decision problem

this situation, the player may wish to start the analysis by considering the particular possible outcomes. Note that two of the outcomes have both the same monetary payoff of \$1000 and the same probability 0.5, given that the respective subtree is chosen. Yet, in general, such outcomes may have different utility values as we mentioned above. For example, in the situation depicted in Figure 11, the prize of \$1000 is given with probability 0.1 in either subtree. However, in the right subtree it comes as a disappointment, whereas in the left subtree it comes as a nice surprise. Returning to Figure 10, denote the utilities by $u(m, s)$ where m is the amount of money and $s \in \{L, R\}$ is the side of the subtree. The normalized utility u of the player may be as follows. $u(-550, L) = 0$, $u(-550, R) = 0.1$, $u(-500, L) = 0.5$, $u(-450, R) = 0.55$, $u(1000, R) = 0.9$, $u(1000, L) = 1$. If so, then the utility value of the left subtree before the outcome is determined is equal to $0.49(0.5) + 0.5(1) = 0.745$, and the utility value of the right subtree is equal to $0.25(0.1)+0.25(0.55)+0.5(0.9) = 0.6125$. It follows that the player prefers the left subtree. The utility value of a randomized choice that picks the left subtree with probability p and the right one with probability $1-p$ should be monotone increasing with p. Note that the set of probability distributions that may arise in this situation consists only of those that are generated by such a randomized choice between the left subtree and the right subtree.

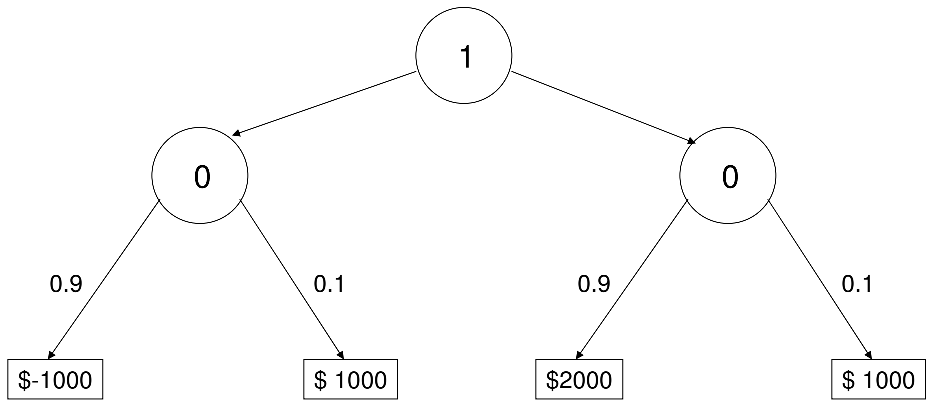

Figure 11: Preferences with regret

Example 3. Consider the situation depicted in Figure 12.

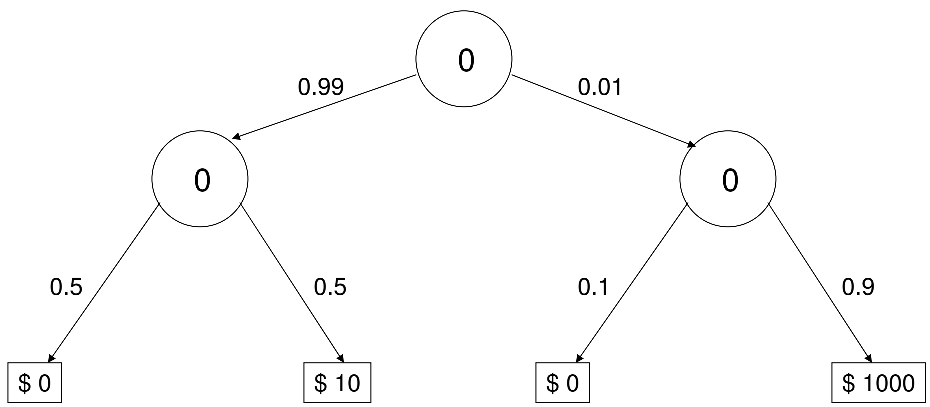

Figure 12: A two-stage lottery

The player may first be surprised if the first lottery picks the right subtree. The left subtree is not as exciting as the right one. In the second stage the player anticipates winning the big prize but may be very disappointed if the loses because the probability of a loss is small. Below we examine one possible way to form utilities in such a situation.

5 Incorporating surprise and regret into utility

Suppose a player receives a prize of m dollars as a result of some random choice. The question is whether the player's utility depends only on the amount m or also on the probability with which the prize was picked, as well as the other prizes and their probabilities. We think that the utility does depend on such factors. For example, if the probability of getting m is only 1%, then the player is pleasantly surprised to get m if the alternative is getting, for example, $m/2$ with probability 99%. Similarly, the player would be very disappointed if the alternative was getting $2m$ with probability 99%. On the other hand, if the probability of getting m is 99%, then there is no big surprise or disappointment if the outcome turns out to be m. However, the actual reaction to an outcome can be more complicated. See Figure 13. Consider, for example,

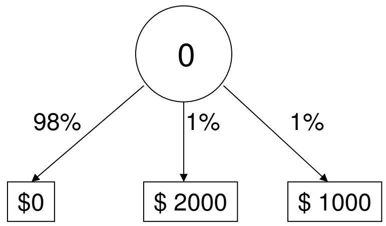

Figure 13: Two similar lotteries with different emotions

a lottery with three outcomes: a prize of \$2000 with probability 1%, a prize of \$1000 with probability 1%, and no prize otherwise. If the outcome is \$1000, then on the one hand the player is quite happy because the prior probability is only 1% and there is an overwhelming probability of no prize. On the other hand, given that an event of probability 1% did occur, the player may be disappointed that he did not get the \$2000. However, if the probability of winning \$2000 was only 0.1% (and zero with probability 98.9%) then the player would not be disappointed with the \$1000 prize, even though there is no significant difference in the distribution and the actual outcome is the same.

Consider first a situation, where there are two possible dollar prizes, m_1 and m_2, with probabilities p_1 and p_2 ($p_1 + p_2 = 1$), respectively. Because there are only two outcomes, a normalized utility function would give one of them the value 1 and the other one the value 0 (unless the player is indifferent between them), so the issues of surprise and disappointment are not reflected by the utility function. However, once a third outcome is present, these issues do get reflected. For simplicity, assume for the moment that there exist two other possible outcomes that are clearly the best possible and the worst possible, respectively, in the whole tree. Thus, these two outcomes have the utility values of 1 and 0, respectively. See Figure 14 Under these

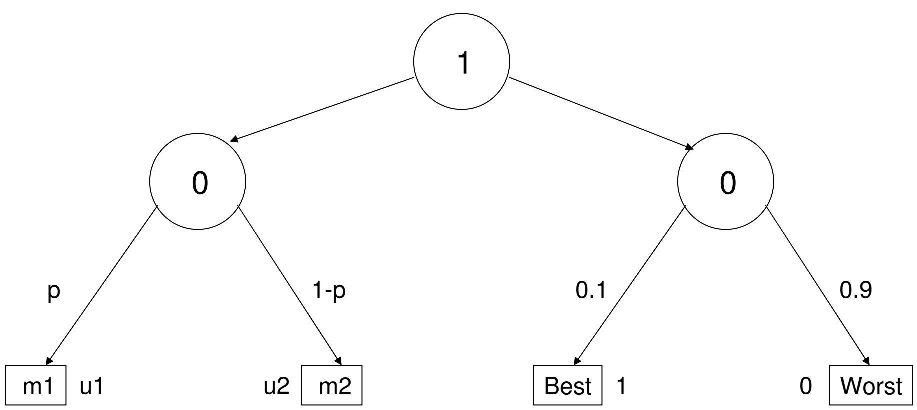

Figure 14: With best and worst

assumptions, with regard to the prizes of m_1 and m_2, we argue that if the outcome m_2 with probability $1-p$ has utility u_2, then the outcome of m_2 with probability p has utility u_1, which depends on u_2, that is given by some function U, i.e.,

$$u_1 = U(u_2) = U(u_2; m_1, p) \ .$$

Similarly, we have u_2 depending on u_1 through the *same* function U,

$$u_2 = U(u_1; m_2, 1-p) \ .$$

Thus,

$$u_1 = U[U(u_1; m_2, 1-p); m_1, p]$$

and

$$u_2 = U[U(u_2; m_1, p); m_2, 1-p] \ ,$$

which means that u_1 is a fixed point of the function

$$f(x) = f(x; m_1, m_2, p) = U[U(x; m_2, 1-p); m_1, p] \ .$$

Similarly, u_2 is a fixed point of the function

$$g(x) = g(x; m_1, m_2, p) = U[U(x; m_1, p); m_2, 1-p] \ .$$

One possible function U can be developed as follows. First, map amounts of money x to $[0, 1]$ by a logistic function

$$\sigma(x) = \frac{1}{1 + \exp\{-(x-\alpha)/\beta\}} \ .$$

Next, use the ratio $\frac{\sigma(m_1)-u_2}{p}$ to measure the satisfaction from m_1 relative to the alternative utility u_2, weighted by the probability of the surprise. Finally, map the ratio back to $[0, 1]$ with the function σ. In summary,

$$U(u_2; m_1, p) = \sigma\left[\frac{\sigma(m_1) - u_2}{p}\right]$$

and

$$U(u_1; m_2, 1-p) = \sigma\left[\frac{\sigma(m_2) - u_1}{1-p}\right] \ .$$

The mapping indicated above is just an example. The actual preferences may require another construction, which we discuss later.

Remark 1. The above idea can extended to more than two outcomes as follows. Suppose the possible outcomes are m_i with probability p_i, $i = 1, \ldots, n$. Suppose the utility from outcome i reflects the enjoyment from m_i with respect to the overall anticipation a priori. Denote the utility from the lottery by u and the utilities from the individual outcomes by u_i. Let us first choose the units so that $\min\{m_i\} = 0$ and $\max\{m_i\} = 1$. On one hand, we may have $u_i = \sigma[(m_i - u)/p_i]$, and on the other hand, $\sum_i p_i u_i = u$. Thus, u can be obtained as the solution of the following equation:

$$f(u) \equiv \sum_{i=1}^{n} p_i \sigma[(m_i - u)/p_i] = u \ . \tag{1}$$

Note that the $f(u)$ is monotone decreasing in terms of u, and for every u, $0 < f(u) < 1$, so the equation (1) has a unique solution in $(0, 1)$.

Remark 2. Suppose a decision maker has to choose one lottery from a set of ℓ possible lotteries. Denote by m_{ij} the dollar amount of the jth prize in the ith lottery ($k_i \geq 1$, $j = 1, \ldots, k_i$, $i = 1, \ldots, \ell$). Denote by p_{ij} the probability with which the jth outcome occurs in the ith lottery ($\sum_{j=1}^{k_i} p_{ij} = 1$; see Figure 15). Note that the prizes are not assumed to be distinct, and there may even be identical pairs $(m_{ij}, p_{ij}) = (m_{hk}, p_{hk})$. Assuming the vNM axioms of utility theory hold, denote by u_{ij} the normalized utility value of the jth

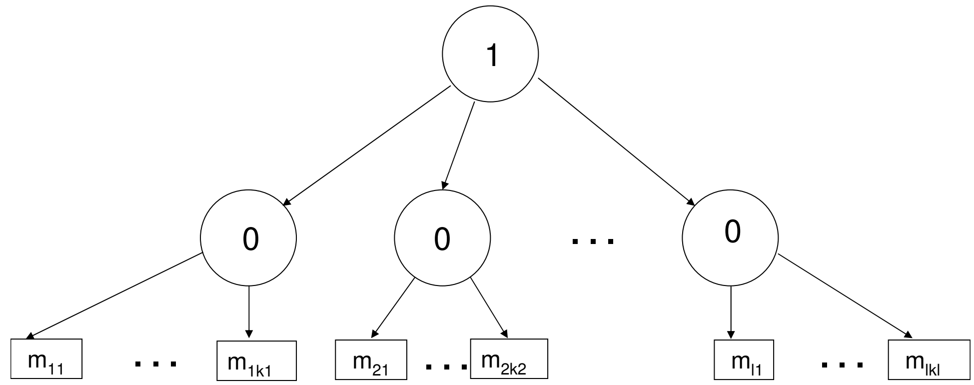

Figure 15: Choosing a lottery

outcome in the ith lottery. In principle, u_{ij} could be a function of all the prizes and all the probabilities in the tree,

$$u_{ij} = u_{ij}((m_{ij}, p_{ij}) : j = 1, \ldots, k_i,\ i = 1, \ldots, \ell)\ .$$

This means that the utility from the ith lottery is equal to $\sum_{j=1}^{k_i} p_{ij}\, u_{ij}$. Note, however, that the theory also gives utility values to every hypothetical lottery over outcomes, but a lottery over the same outcomes of the ith lottery but with different probabilities is counterfactual, because the specification of the outcome includes the probability with which the outcome occurs. If the axioms are relaxed required to hold only over lotteries that are not counterfactual, then uniqueness of utility may not hold.

6 Money-equivalent values

Representing preferences with utility functions may be possible theoretically, but in practice is not easy for a decision maker to implement. However, people may find it easier to think in more practical terms and measure their utilities in terms of money. More concretely, one possibility to represent preferences over probability distributions over outcomes is to estimate the "money-equivalent" value. Given an option to participate in a certain lottery, the decision maker can ask himself what a sure monetary prize should be so that he would accept it instead of the lottery. For example, suppose the decision maker holds a lottery ticket that gives \$1000 with probability p and zero otherwise. Taking into account all factors, such as surprise and regret, the decision maker should be able to specify a minimum amount of money $m = m(p)$ for which he would be wiling to sell the ticket. Among other factors, the value $m(p)$ may also depend on whether or not the decision maker would later be able to find out the outcome of the lottery, regret selling the ticket if he would have won, or be happy that he sold the ticket in case it would have lost. The function $m(p)$ would likely be monotone in p, but we do not make any assumptions about it at this point.

The function $m(p)$ could be related to vNM utility as follows. Assume for a moment that the preferences of the decision maker (over probability distribution over monetary prizes) satisfy the utility axioms of von Neumann and Morgenstern. In this case, the player has a utility function $u(\cdot)$ over monetary prizes, which extends by expectation to probability distributions. If the utility satisfies $u(\$1000) = 1$ and $u(0) = 0$, then $u[p \cdot \$1000 + (1 - p) \cdot \$0] = p$. If the decision maker is indifferent between receiving for sure an amount M of money and holding the lottery ticket, then $u(M) = p$. Thus, $u(M) = u(m(p)) = p$, hence $m(p) = u^{-1}(p)$, i.e., as functions $u = m^{-1}$. However, in general, assigning monetary values to outcomes is not equivalent to having a vNM utility for money, because an outcome entails more than just the amount of money received, as we discussed above. The money-equivalent value reflects all the circumstances related to the point at which it is applied.

With regard to the prospect theory, if the decision maker has a value function of the form $v(p) = \pi(p)V(\$1000) + \pi(1-p)V(0)$ and $V(0) = 0$, then $v(p) = \pi(p)V(\$1000)$. If the function $V(\cdot)$ extends monotonically and continuously to all sure monetary prizes, then by definition $V(m(p)) = \pi(p)\,V(\$1000)$, hence $m(p) = V^{-1}[\pi(p)\,V(\$1000]$.

The question now arises whether or not there is any methodology for deriving money-equivalent values. We propose one example below.

As we discussed above, the satisfaction from an outcome in which the player receives a certain amount of money may depend may depend on the prior probability of the outcome, as well as on the anticipation of the player from the lottery or the whole situation. Consider a lottery that gives m_i dollars with probability p_i, $i = 1, \ldots, k$. Denote by u the money-equivalent value of the lottery. We use u as a reference level for determining the money-equivalent values of the individual outcomes. Thus, we rely on some function $v(m, u, p)$, so that $u_i = v(m_i, u, p_i)$, where

(i) $v(m, m, p) = m$,

(ii) $v(m, u, p) > m$ if $u < m$ and $v(m, u, p) < m$ if $m < u$,

(iii) $v(m, u, p)$ decreases with p if $m > u$ and increases with p if $m < u$, and

(iv) $v(m, u, p)$ is monotone decreasing in terms of u.

One possibility is the following:

$$v(m, u, p) = m \cdot \left[1 + \frac{\alpha(m-u)}{(1+p)\max\{m, u\}}\right] .$$

Note that if $m > u$, then

$$v(m, u, p) = m \cdot \left[1 + \frac{\alpha(m-u)}{(1+p)m}\right] ,$$

and if $m < u$) then

$$v(m, u, p) = m \cdot \left[1 + \frac{\alpha(m-u)}{(1+p)u}\right] = m \cdot \left[1 + \frac{\alpha(m/u-1)}{1+p}\right] .$$

Thus, $v(m, u, p)$ is monotone decreasing in terms of u. Next, let $U(m)$ be a vNM utility function in terms of a sure monetary prize. We assume $U(0) = 0$. By definition, the following equation must hold:

$$U(u) = \sum_{i=1}^{k} p_i U[v(m_i, u, p_i)] . \tag{2}$$

The value of u is therefore determined from the latter. Note that the left-hand side is monotone increasing in terms of u and the right-hand side is monotone decreasing in terms of u. Once u has been determined, the money-equivalent of the ith outcome can be evaluated as $v(m_i, u, p_i)$.

The fixed-point property of u can be explained as follows. Consider any tentative value x for the money-equivalent value of the lottery. If indeed this is correct money-equivalent value of the lottery, then the money-equivalent value of outcome i should be equal to $v(m_i, x, p_i)$. The utility from it should be equal to $U[v(m_i, x, p_i)]$, hence the expected utility from the lottery should be equal to $\sum_{i=1}^{k} p_i U[v(m_i, x, p_i)]$. Hence, the money-equivalent value of the lottery should be equally to $x' = U^{-1}\left(\sum_{i=1}^{k} p_i U[v(m_i, x, p_i)]\right)$. Thus, only at a fixed point defined in (2) this money-equivalent derived value x' is equal to the value x it was assumed to have.

In principle, the money-equivalent value u_i of an outcome may depend on money than just on m_i, p_i and the money-equivalent value of the whole lottery. It may of the form

$$u_i = V_i(m_i, p_i; u_j : j \neq i) \ .$$

so, the fixed point is in a higher dimension.

7 Conditional utilities

The situation can be much more complicated if instead of a lottery built into a game tree, the choices are made by a random device that implements a mixed strategy of the player himself or of another player. Thus, no prior probabilities are associated with the outcomes. The player for his preferences while considering choosing a mixed strategy or imagining mixed strategies of other players. In equilibrium, the players must have analyzed the situation and formed their preferences with respect to each other's choice of mixed strategies. When Player 1 considers possible outcomes, he may fix in his mind probabilities with which Player 2 picks the outcomes. The probabilities can constitute Player 1's beliefs, but they can also form a specific mixed strategy of some other player, which can be part of an equilibrium. If, indeed, the preferences of Player 1 over outcomes depend on the probabilities with which they occur, then a utility function cannot be fixed a priori, unless we assume a particular function that related the utility to such probabilities. Thus, in the situation depicted in Figure 15, the player can consider any strategy that picks lottery i with probability x_i, $i = 1, \ldots, \ell$. Such a mixed strategy induces a probability distribution on the final outcomes, where the jth outcome of lottery i occurs with probability $x_i \, p_{ij}$, but we do not wish to assume that the player is indifferent between that distribution and playing the mixed strategy $(x_1, \ldots, x_n)$. Instead, we wish to develop values u_{ij} as *conditional utilities* of outcome j, given that the ith lottery has been chosen. This purpose of this notion of conditional utility is to separate the surprise of the ith lottery having been chosen from the values of the various outcomes it entails. Furthermore, the value u_i derived above as a fixed point, can also be interpreted as a conditional utility of the lottery i, given that it was chosen. Again, the purpose of this concept is to separate the surprise or regret of the ith lottery having been chosen from the aggregate anticipation of the values of the various outcomes it entails. Given these values, we may derive the *unconditional* utility U_i of reaching the point where the ith lottery is about to be performed as follows. Given the mixed strategy $(x_1, \ldots, x_n)$, we seek a fixed point u^* of the equation

$$U(u^*) = \sum_{i=1}^{\ell} x_i U[v(u_i, u^*, x_i)] \ ,$$

and set

$$U_i = U[v(u_i, u^*, x_i)] \ .$$

Thus,

$$u^* = \sum_i x_i \, U_i \ .$$

We believe that the concepts of conditional utility and money-equivalent value may be the right ones for analyzing a in extensive form. Conditional utility isolates a subtree from the rest of the game so that players are not concerned with comparisons to what could have been the outcome if that subtree was not chosen. However, this concept still does not address the problem utility that depends on attitudes toward other players in view of their past choices as discussed in [5]. This issue still presents a challenge to the theory as we demonstrate in the example below.

In the game depicted in Figure 16 there is one prize of \$1M. The probability with which the prize is awarded to one of the players depends on the choices the player make. With high probability, the prize is not awarded at all. If Player I choose Left in the first step, then he gets the prize with probability 0.01%. If he chooses Right, the player has can choose Left and get the prize with probability 0.1%. The probability of winning the prize grows exponentially if a player decides not to terminate the game, each time the other player has the chance to win. In the last move, Player II choosing between taking the full 10% to himself or sharing it with Player I. A deterministic version of this game was first proposed in [4] and experiments on it were reported in [3]. The prospect-theory approach proposed by Kahneman and Tversky assigns values V_0 and V_1 to losing and winning the prize, respectively, and uses a monotone increasing weighting function $\pi(p)$ so that the value of a lottery in which a player wins the prize with probability p is equal to $\pi(p)V_1+\pi(1-p)V_0$. It follows that both players have such utilities in this game, then the only equilibrium point is where Player 1 chooses Left and terminates the game. The same problem arises with any approach that uses only the monetary values and monotone increasing functions of the probability of winning the prize.

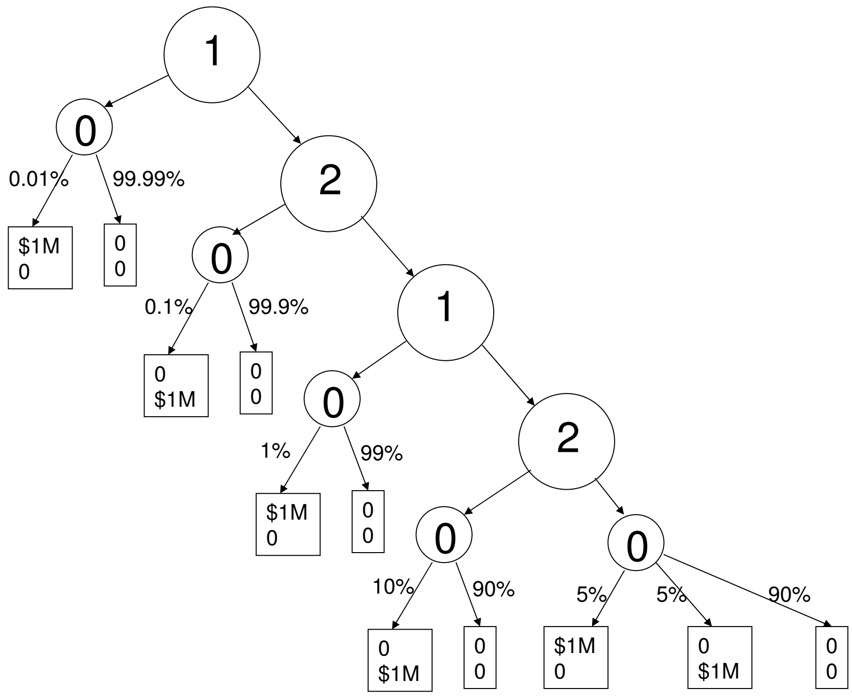

Figure 16: An example with two monetary prizes

The main problem here is that a money-equivalent value of a subgame has to be determined with respect to the past decisions that have led into the subtree. The subgame itself does not have that information. Possible guilt feelings of the player who terminates the game have to be taken into account.

8 Games in normal form

When a game is played normal form there are only two relevant time points: decision time and reward time. In view of our discussion above, the normal of a game given in extensive (tree) form cannot be equivalent to the extensive form. We are still concerned though with the dependence of utility of outcomes on surprise and regret. The latter can be the result of coin tosses implementing mixed strategies.

Consider the two-person zero sum game in normal form in Figure 17. One representation of this game in extensive form is shown in Figure 18. Suppose the row player plays the mixed strategy $(x, 1-x)$ and the column player plays the mixed strategy $(y, 1-y)$. There are various ways to interpret the normal form

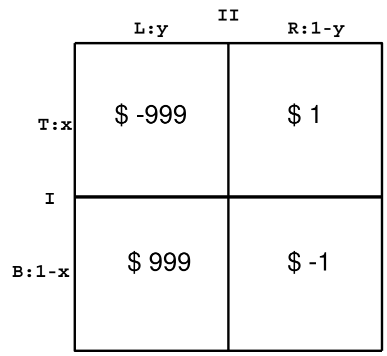

Figure 17: A two-Person zero-Sum game with small probabilities in equilibrium

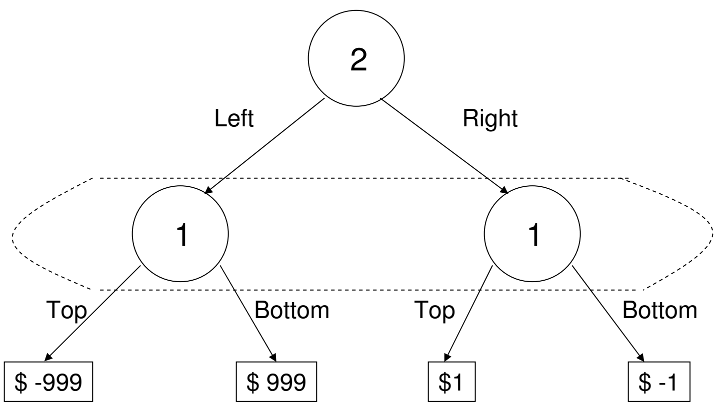

Figure 18: A two-Person zero-Sum game with small probabilities in equilibrium

of a game. Suppose the players themselves implement their own mixed strategies rather than tell a referee to toss their coins themselves. Thus, each player knows the outcome of his own toss before he knows the outcome of the his opponents' coin toss. In this setting, the two player may have quite different experiences. Suppose the players consider using the monetary amount as vNM utilities. The optimal strategies in this case are $x = 1/2$ and $y = 1/1000$. Thus, the row player is not surprised with any outcome of his coin. The column player though believes a priori that he will have to play R, and is very surprised if he has to play L, which is more risky for him. It is interesting to mention in this regard that the outcome of one's own coin toss is not likely to be a nice surprise because he had the option to act according to that outcome anyway. However, is prepared play a certain mixed strategy and then the coin tells him to play some strategy that has only a small probability a priori, then he may be afraid to follow the coin if the strategy seems too risky.

If a pair (x, y) of mixed strategies is in equilibrium with respect to the true utilities, then nothing changes if the players reveal their mixed strategies to each other. Now, suppose the row player believes the column player plays $(y, 1 - y)$. Then, his own choice is between a bet T in which he loses \$999 with probability y and wins \$1 with probability $1 - y$, and a bet B in which he wins \$999 with probability y and loses \$1 with probability $1 - y$. See Figure 19. If $y = 0$, then the row player wins \$1 if he chooses T and loses one dollar if he choose B, so he prefers T. If $y = 1$, then by a similar argument he prefers B. The vNM axioms imply that there exists a y at which the row player is indifferent between T an B, and therefore also between them and any lottery between them. It follows, that *in equilibrium, a player should not have any strong feelings*

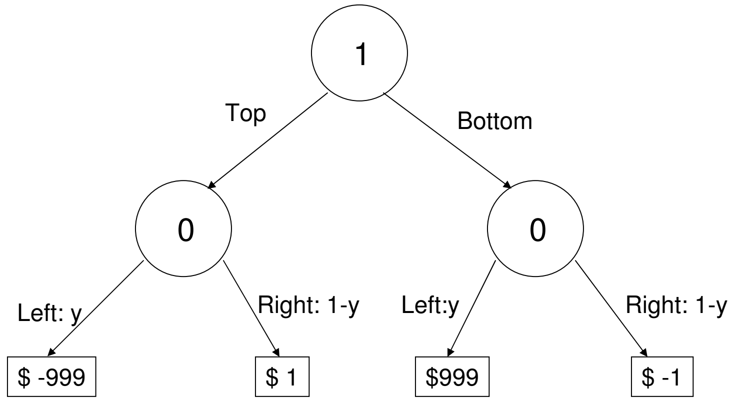

Figure 19: The bet with resect to opponent's coin toss

or emotions after he has learned the outcome of his own coin toss but before he has learned the outcome of his opponent's coin toss, because he is indifferent between the alternatives, i.e., the two bets induced by y. However, once the outcome of opponent's coin toss is disclosed, the row player may have very strong feelings about the choice he made after his own coin toss.

References

[1] J. Anderoni and J. H. Miller, "Rational cooperation in the finitely repeated Prisoner's Dilemma: Experimental evidence," *The Economic Journal* **103** (1993) 570–585.

[2] D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* **47** (1979) 263–292.

[3] R. D. McKelvey and T. R. Palfrey, "An Experimental Study of the Centipede Game," *Econometrica* **60** (1992) 803–836.

[4] N. Megiddo, "Remarks on Bounded Rationality," Research Report RJ 5270, IBM Almaden Research Center, San Jose, California, 1986.
`http://theory.stanford.edu/ megiddo/pdf/RemarksRJ5270.pdf`

[5] N. Megiddo, "Formation of Preferences and Strategic Analysis: Can they be De-Coupled?" 9/24/2003. Presented at IMGTA'2001- XIV Italian Meeting on Game Theory and Applications.
`http://theory.stanford.edu/∼megiddo/pdf/ver3.092403.pdf`

[6] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Third Edition, Princeton University Press, 1953.