A Model of Casino Gambling

Nicholas Barberis

Yale University

June 2011*

Abstract

We show that prospect theory offers a rich theory of casino gambling, one that captures several features of actual gambling behavior. First, we demonstrate that, for a wide range of preference parameter values, a prospect theory agent would be willing to gamble in a casino even if the casino only offers bets with no skewness and with zero or negative expected value. Second, we show that the probability weighting embedded in prospect theory leads to a plausible time inconsistency: at the moment he enters a casino, the agent plans to follow one particular gambling strategy; but after he starts playing, he wants to switch to a different strategy. The model therefore predicts heterogeneity in gambling behavior: how a gambler behaves depends on whether he is aware of the time inconsistency; and, if he is aware of it, on whether he can commit in advance to his initial plan of action.

JEL classification: D03

Keywords: gambling, prospect theory, time inconsistency, probability weighting

*I am grateful to Isaiah Andrews and Ji Shen for excellent research assistance. I am also grateful, for very helpful comments, to Ulf Axelson, Daniel Benjamin, Stefano DellaVigna, Yaxin Duan, Keith Gamble, Daniel Gottlieb, Jonathan Ingersoll, Bjorn Johnson, Kristof Madarasz, Ted O'Donoghue, Matthew Rabin, Georg Weizsacker, and seminar participants at Caltech, Columbia University, Cornell University, DePaul University, Harvard University, the London School of Economics, Princeton University, Stanford University, the Stockholm School of Economics, the University of California at Berkeley, the University of Chicago, the Wharton School, Yale University, and the NBER.

1 Introduction

Casino gambling is a hugely popular activity. The American Gaming Association reports that, in 2007, 55 million people made 376 million trips to casinos in the United States alone.

If we are to fully understand how people think about risk, we need to make sense of the existence and popularity of casino gambling. Unfortunately, there are still very few models of why people go to casinos or of how they behave when they get there. The challenge is clear. In the field of economics, the standard model of decision-making under risk couples the expected utility framework with a concave utility function defined over wealth. This model is helpful for understanding a range of phenomena. It cannot, however, explain casino gambling: an agent with a concave utility function will always turn down a wealth bet with a negative expected value.

While casino gambling is not consistent with the standard economic model of risk attitudes, researchers have made some progress in understanding it better. One approach is to introduce non-concave segments into the utility function (Friedman and Savage, 1948). A second approach argues that people derive a separate component of utility from gambling. This utility may be only indirectly related to the bets themselves – for example, it may stem from the social pleasure of going to a casino with friends – or it may be directly related to the bets, in that the gambler enjoys the feeling of suspense as he waits for the bets to play out (Conlisk, 1993). A third approach suggests that gamblers are simply unaware that the odds they are facing are unfavorable.

In this paper, we present a new model of casino gambling based on Tversky and Kahneman's (1992) cumulative prospect theory. Cumulative prospect theory, a prominent theory of decision-making under risk, is a modified version of Kahneman and Tversky's (1979) prospect theory. It posits that people evaluate risk using a value function that is defined over gains and losses, that is concave over gains and convex over losses, and that is kinked at the origin, so that people are more sensitive to losses than to gains, a feature known as loss aversion. It also states that people engage in "probability weighting": that they use *transformed* rather than objective probabilities, where the transformed probabilities are obtained from objective probabilities by applying a weighting function. The main effect of the weighting function is to overweight the tails of the distribution it is applied to. The overweighting of tails does not represent a bias in beliefs; rather, it is a way of capturing the common preference for a lottery-like, or positively skewed, payoff.[1]

[1] Although our model is based on the cumulative prospect theory of Tversky and Kahneman (1992) rather than on the original prospect theory of Kahneman and Tversky (1979), we will sometimes refer to cumulative

We choose prospect theory as the basis for a possible explanation of casino gambling because we would like to understand gambling in a framework that also explains other evidence on risk attitudes. Prospect theory can explain a wide range of experimental evidence on attitudes to risk – indeed, it was designed to – and it can also shed light on much field evidence on risk-taking: for example, it can address a number of facts about risk premia in asset markets (Benartzi and Thaler, 1995; Barberis and Huang, 2008). By offering a prospect theory model of casino gambling, our paper therefore suggests that gambling is not an isolated phenomenon requiring its own unique explanation, but that it may instead be one of a family of facts that can be understood using a single model of risk attitudes.

The idea that prospect theory might explain casino gambling is initially surprising. Through the overweighting of the tails of distributions, prospect theory can easily explain why people buy lottery tickets. However, many casino games offer gambles that, aside from their low expected values, are also much less skewed than a lottery ticket. It is therefore far from clear that probability weighting can explain why these gambles are so popular. Indeed, given that prospect theory agents are much more sensitive to losses than to gains, one would think that they would find these gambles very unappealing. Initially, then, prospect theory does not seem to be a promising starting point for a model of casino gambling.

In this paper, we show that, in fact, prospect theory can offer a rich theory of casino gambling, one that captures several features of actual gambling behavior. First, we demonstrate that, for a wide range of preference parameter values, a prospect theory agent *would* be willing to gamble in a casino, even if the casino only offers bets with no skewness and with zero or negative expected value. Second, we show that prospect theory – in particular, its probability weighting feature – predicts a plausible *time inconsistency*: at the moment he enters a casino, a prospect theory agent plans to follow one particular gambling strategy; but after he starts playing, he wants to switch to a different strategy. How he behaves therefore depends on whether he is aware of the time inconsistency; and, if he *is* aware of it, on whether he is able to commit in advance to his initial plan of action.

What is the intuition for why, in spite of his loss aversion, a prospect theory agent might still be willing to enter a casino? Consider a casino that offers only zero expected value bets – specifically, 50:50 bets to win or lose some fixed amount $\$h$ – and suppose that the agent makes decisions by maximizing the cumulative prospect theory value of his accumulated winnings or losses at the moment he leaves the casino. We show that, if the agent enters the casino, his preferred plan is usually to keep gambling if he is winning, but to stop gambling

prospect theory as "prospect theory" for short.

and leave the casino if he starts accumulating losses. An important property of this plan is that, even though the casino offers only 50:50 bets, the distribution of the agent's perceived *overall* casino winnings becomes positively skewed: by stopping once he starts accumulating losses, the agent limits his downside; and by continuing to gamble when he is winning, he retains substantial upside.

At this point, the probability weighting feature of prospect theory plays an important role. Under probability weighting, the agent overweights the tails of probability distributions. With sufficient probability weighting, then, the agent may *like* the positively skewed distribution generated by his planned gambling strategy. We show that, for a wide range of preference parameter values, the probability weighting effect indeed outweighs the loss aversion effect and the agent *is* willing to enter the casino. In other words, while the prospect theory agent would always turn down the basic 50:50 bet if it were offered *in isolation*, he is nonetheless willing to enter the casino because, through a specific choice of exit strategy, he gives his overall casino experience a positively skewed distribution, one which, with sufficient probability weighting, he finds attractive.

Prospect theory offers more than just an explanation of why people go to casinos; it also predicts a time inconsistency. The inconsistency is a consequence of probability weighting: it arises because, as time passes, the probabilities of final outcomes change, which, in turn, means that the degree to which the agent under- or overweights these outcomes also changes. For example, when he enters the casino, the agent knows that the probability of winning five bets in a row, and hence of accumulating a total of $\$5h$, is very low, namely $\frac{1}{32}$. Under probability weighting, a low probability outcome like this is overweighted. If the agent actually wins the first four bets, however, the probability of winning the fifth bet, and hence of accumulating $\$5h$, is now $\frac{1}{2}$. Under probability weighting, a moderate probability outcome like this is *under*-weighted.

The fact that some final outcomes are initially overweighted but subsequently underweighted, or vice-versa, means that the agent's preferences over gambling strategies change over time. We noted above, that, at the moment he enters a casino, the agent's preferred plan is usually to keep gambling if he is winning but to stop gambling if he starts accumulating losses. We show, however, that once he starts playing, he wants to do the opposite: to keep gambling if he is losing and to stop if he accumulates a significant gain.

As a result of this time inconsistency, our model predicts heterogeneity in gambling behavior. How a gambler behaves depends on whether he is aware of the time inconsistency. A gambler who *is* aware of the time inconsistency has an incentive to try to commit to his initial plan of action. For gamblers who are aware of the time inconsistency, then, their behavior further depends on whether they are indeed able to find a commitment device.

To study these distinctions, we consider three types of agents. The first type is "naive": he is unaware of the time inconsistency. This agent typically *plans* to keep gambling if he is winning and to stop if he starts accumulating losses. After he starts playing, however, he deviates from this plan and instead gambles as long as possible when he is losing and stops if he accumulates a significant gain.

The second type of agent is "sophisticated" – he is aware of the time inconsistency – but is unable to find a way of committing to his initial plan. He therefore knows that, if he enters the casino, he will keep gambling if he is losing and will stop if he makes some gains. This will give his overall casino experience a *negatively* skewed distribution. Since he overweights the tails of probability distributions, he almost always finds this unattractive and therefore refuses to enter the casino in the first place.

The third type of agent is also sophisticated but is able to find a way of committing to his initial plan. Just like the naive agent, this agent typically plans, on entering the casino, to keep gambling if he is winning and to stop if he starts accumulating losses. Unlike the naive agent, however, he is able, through the use of a commitment device, to stick to this plan. For example, he may bring only a small amount of cash to the casino while also leaving his ATM card at home; this guarantees that he will indeed leave the casino if he starts accumulating losses.

In summary, under the view proposed in this paper, the popularity of casinos is driven by two aspects of our psychological makeup: first, by the tendency to overweight the tails of distributions, which makes even the small chance of a large win seem very alluring; and second, by what we could call "naivete," namely the failure to recognize that, after starting to gamble, we may deviate from our initial plan of action.

Our model is a complement to existing theories of gambling, not a replacement. For example, we suspect that the concept of "utility of gambling" plays at least as large a role in casinos as does prospect theory. At the same time, we think that prospect theory can add significantly to our understanding of casino gambling. As noted above, one attractive feature of the prospect theory approach is that it not only explains why people go to casinos, but also offers a rich description of what they do once they get there. In particular, it explains a number of features of casino gambling that have not emerged from earlier models: the tendency to gamble longer than planned when losing, the strategy of leaving one's ATM card at home, and casinos' practice of issuing vouchers for free food and accommodation to people who are winning.[2]

In recent years, there has been a surge of interest in the time inconsistency that stems from hyperbolic discounting. In this paper, we study a different time inconsistency, one generated by probability weighting. The prior literature offers very little guidance on how best to analyze this particular inconsistency. We present an approach that we think is simple and natural; but other approaches are certainly possible.

Our focus is on the "demand" side of casino gambling: we posit a casino structure and study a prospect theory agent's reaction to it. In Section 4.2, we briefly discuss an analysis of the "supply" side – of what kinds of gambles we should expect to see offered in an economy with prospect theory agents. However, we defer a full analysis of the supply side to future research.[3]

2 Cumulative Prospect Theory

In this section, we review the elements of cumulative prospect theory. Readers who are already familiar with this theory may prefer to go directly to Section 3.

Consider the gamble

$$(x_{-m}, p_{-m}; \ldots; x_{-1}, p_{-1}; x_0, p_0; x_1, p_1; \ldots; x_n, p_n), \tag{1}$$

to be read as "gain x_{-m} with probability p_{-m}, x_{-m+1} with probability p_{-m+1}, and so on, independent of other risks," where $x_i < x_j$ for $i < j$, $x_0 = 0$, and $\sum_{i=-m}^{n} p_i = 1$. In the expected utility framework, an agent with utility function $U(\cdot)$ evaluates this gamble by computing

$$\sum_{i=-m}^{n} p_i U(W + x_i), \tag{2}$$

where W is his current wealth. A cumulative prospect theory agent, by contrast, assigns the gamble the value

$$\sum_{i=-m}^{n} \pi_i v(x_i), \tag{3}$$

[2]It would be interesting to incorporate an explicit utility of gambling into the model we present below. The only reason we do not do so is because we want to understand the predictions of prospect theory, taken alone.

[3]A sizeable literature in medical science studies "pathological gambling," a disorder which affects about 1% of gamblers. Our paper is not aimed at understanding such extreme gambling behavior, but rather the behavior of the vast majority of gamblers whose activity is not considered medically problematic.

where[4]

$$\pi_i = \begin{cases} w(p_i + \ldots + p_n) - w(p_{i+1} + \ldots + p_n) & \text{for} \quad 0 \leq i \leq n \\ w(p_{-m} + \ldots + p_i) - w(p_{-m} + \ldots + p_{i-1}) & \phantom{\text{for}} \quad -m \leq i < 0 \end{cases}, \tag{4}$$

and where $v(\cdot)$ and $w(\cdot)$ are known as the value function and the probability weighting function, respectively. Tversky and Kahneman (1992) propose the functional forms

$$v(x) = \begin{cases} x^\alpha & \text{for} \quad x \geq 0 \\ -\lambda(-x)^\alpha & \phantom{\text{for}} \quad x < 0 \end{cases} \tag{5}$$

and

$$w(P) = \frac{P^\delta}{(P^\delta + (1-P)^\delta)^{1/\delta}}, \tag{6}$$

where α, $\delta \in (0,1)$ and $\lambda > 1$. The left panel in Figure 1 plots the value function in (5) for $\alpha = 0.5$ and $\lambda = 2.5$. The right panel in the figure plots the weighting function in (6) for $\delta = 0.4$ (the dashed line), for $\delta = 0.65$ (the solid line), and for $\delta = 1$, which corresponds to no probability weighting at all (the dotted line). Note that $v(0) = 0$, $w(0) = 0$, and $w(1) = 1$.

There are four important differences between (2) and (3). First, the carriers of value in cumulative prospect theory are gains and losses, not final wealth levels: the argument of $v(\cdot)$ in (3) is x_i, not $W + x_i$. Second, while $U(\cdot)$ is typically concave everywhere, $v(\cdot)$ is concave only over gains; over losses, it is convex. This captures the experimental finding that people tend to be risk averse over moderate-probability gains – they prefer a certain gain of \$500 to $(\$1000, \frac{1}{2})$ – but risk-*seeking* over moderate-probability losses, in that they prefer $(-\$1000, \frac{1}{2})$ to a certain loss of \$500.[5] The degree of concavity over gains and of convexity over losses are both governed by the parameter α; a lower value of α means greater concavity over gains and greater convexity over losses. Using experimental data, Tversky and Kahneman (1992) estimate $\alpha = 0.88$ for their median subject.

Third, while $U(\cdot)$ is typically differentiable everywhere, the value function $v(\cdot)$ is kinked at the origin so that the agent is more sensitive to losses – even small losses – than to gains of the same magnitude. As noted in the Introduction, this element of cumulative prospect theory is known as loss aversion and is designed to capture the widespread aversion to bets such as $(\$110, \frac{1}{2}; -\$100, \frac{1}{2})$. The severity of the kink is determined by the parameter λ; a higher value of λ implies a greater relative sensitivity to losses. Tversky and Kahneman (1992) estimate $\lambda = 2.25$ for their median subject.

Finally, under cumulative prospect theory, the agent does not use objective probabilities

[4]When $i = n$ and $i = -m$, equation (4) reduces to $\pi_n = w(p_n)$ and $\pi_{-m} = w(p_{-m})$, respectively.

[5]We abbreviate $(x, p; 0, q)$ to (x, p).

when evaluating a gamble, but rather, transformed probabilities obtained from objective probabilities via the weighting function $w(\cdot)$. Equation (4) shows that, to obtain the probability weight π_i for an outcome $x_i \geq 0$, we take the total probability of all outcomes equal to or better than x_i, namely $p_i + \ldots + p_n$, the total probability of all outcomes strictly better than x_i, namely $p_{i+1} + \ldots + p_n$, apply the weighting function to each, and compute the difference. To obtain the probability weight for an outcome $x_i < 0$, we take the total probability of all outcomes equal to or worse than x_i, the total probability of all outcomes strictly worse than x_i, apply the weighting function to each, and compute the difference.[6]

The main consequence of the probability weighting in (4) and (6) is that the agent overweights the *tails* of any distribution he faces. In equations (3)-(4), the most extreme outcomes, x_{-m} and x_n, are assigned the probability weights $w(p_{-m})$ and $w(p_n)$, respectively. For the functional form in (6) and for $\delta \in (0,1)$, $w(P) > P$ for low, positive P; the right panel of Figure 1 illustrates this for $\delta = 0.4$ and $\delta = 0.65$. If p_{-m} and p_n are small, then, we have $w(p_{-m}) > p_{-m}$ and $w(p_n) > p_n$, so that the most extreme outcomes – the outcomes in the tails – are overweighted.

The overweighting of tails in (4) and (6) is designed to capture the simultaneous demand many people have for both lotteries and insurance. For example, subjects typically prefer $(\$5000, 0.001)$ to a certain \$5, but also prefer a certain loss of \$5 to $(-\$5000, 0.001)$. By overweighting the tail probability of 0.001 sufficiently, cumulative prospect theory can capture both of these choices. The degree to which the agent overweights tails is governed by the parameter δ; a lower value of δ implies more overweighting of tails. Tversky and Kahneman (1992) estimate $\delta = 0.65$ for their median subject. To ensure the monotonicity of $w(\cdot)$, we require $\delta \in (0.28, 1)$.[7]

We emphasize that the transformed probabilities in (3)-(4) do not represent erroneous beliefs: in Tversky and Kahneman's framework, an agent evaluating the lottery-like $(\$5000, 0.001)$ gamble knows that the probability of receiving the \$5000 is exactly 0.001. Rather, the transformed probabilities are decision weights that capture the experimental evidence on risk attitudes – for example, the preference for the lottery over a certain \$5.

[6] The main difference between cumulative prospect theory and the original prospect theory in Kahneman and Tversky (1979) is that, in the original version, the weighting function $w(\cdot)$ is applied to the probability density function rather than to the cumulative probability distribution. By applying the weighting function to the cumulative distribution, Tversky and Kahneman (1992) ensure that cumulative prospect theory satisfies the first-order stochastic dominance property. The original prospect theory, by contrast, does not satisfy this property.

[7] To be precise, Tversky and Kahneman (1992) allow the value of δ to depend on whether the outcome that is being assigned a probability weight is a gain or a loss. They estimate $\delta = 0.61$ for gain outcomes and $\delta = 0.69$ for loss outcomes. For simplicity, we use the same value of δ for both gain and loss outcomes and take its median estimate to be the average of 0.61 and 0.69, namely 0.65.

3 A Model of Casino Gambling

In the United States, the term "gambling" typically refers to one of four things: (i) casino gambling, of which the most popular forms are slot machines and the card game of blackjack; (ii) the buying of lottery tickets; (iii) pari-mutuel betting on horses at racetracks; and (iv) fixed-odds betting through bookmakers on sports such as football, baseball, basketball, and hockey. The American Gaming Association estimates the 2007 revenues from the four types of gambling at $60 billion, $24 billion, $4 billion, and $200 million, respectively.[8]

While the four types of gambling listed above have some common characteristics, they also differ in some ways. Casino gambling differs from playing the lottery in that casino games offer bets that are typically much less positively skewed than a lottery. And it differs from racetrack-betting and sports-betting in that casino games usually require less skill.

In this paper, we focus on casino gambling, largely because, from the perspective of prospect theory, it seems particularly hard to explain. The buying of lottery tickets is already captured by prospect theory through the overweighting of tail probabilities. However, since casino games offer bets that are much less positively skewed than lotteries, it is not at all clear that we can use the overweighting of tails to explain their popularity.

We model a casino in the following way. There are $T+1$ dates, $t = 0, 1, \ldots, T$. At time 0, the casino offers the agent a 50:50 bet to win or lose a fixed amount $\$h$. If the agent turns the gamble down, the game is over: he is offered no more gambles and we say that he has declined to enter the casino. If the agent *accepts* the 50:50 bet, we say that he has agreed to enter the casino. The gamble is then played out and, at time 1, the outcome is announced. At that time, the casino offers the agent another 50:50 bet to win or lose $\$h$. If he turns it down, the game is over: the agent settles his account and leaves the casino. If he *accepts* the gamble, it is played out and, at time 2, the outcome is announced. The game then continues in the same way. If, at time $t \in [0, T-2]$, the agent agrees to play a 50:50 bet to win or lose $\$h$, then, at time $t+1$, he is offered another such bet and must either accept it or decline it. If he declines it, the game is over: he settles his account and leaves the casino. At time T, the agent *must* leave the casino if he has not already done so. We think of the interval from 0 to T as an evening of play.

We assume that there is an exogeneous date, date T, at which the agent must leave the casino if he has not already done so because we think that this makes the model more

[8] The $200 million figure corresponds to sports-betting through *legal* bookmakers. It is widely believed that this figure is dwarfed by the revenues from illegal sports-betting. Also excluded from these figures are the revenues from online gambling.

realistic: whether because of fatigue or because of work and family commitments, most people cannot stay in a casino indefinitely.[9]

Of the major casino games, it is blackjack that most closely matches the game in our model: for a player familiar with the basic strategy, the odds of winning a round of blackjack are close to 0.5. Slot machines offer a positively skewed payoff and therefore, at first sight, do not appear to fit the model as neatly. Later, however, we argue that our analysis may be able to shed as much light on slot machines as it does on blackjack.

In the discussion that follows, it will be helpful to think of the casino as a binomial tree. Figure 2 illustrates this for $T = 5$ – ignore the arrows, for now. Each column of nodes in the tree corresponds to a particular time: the left-most node corresponds to time 0 and the right-most column to time T. At time 0, then, the agent starts in the left-most node. If he takes the time 0 bet and wins, he moves one step up and to the right; if he takes the time 0 bet and loses, he moves one step down and to the right, and so on: whenever the agent wins a bet, he moves up a step in the tree, and whenever he loses, he moves down a step. The various nodes in any given column therefore represent the different possible accumulated winnings or losses at that time.

We refer to each node in the tree by a pair of numbers (t, j). The first number, t, which ranges from 0 to T, indicates the time that the node corresponds to. The second number, j, which, for given t, ranges from 1 to $t + 1$, indicates how far down the node is within the column of nodes that corresponds to time t: the highest node in the column corresponds to $j = 1$ and the lowest node to $j = t + 1$. The left-most node in the tree is therefore node $(0, 1)$. The two nodes in the column immediately to the right, starting from the top, are nodes $(1, 1)$ and $(1, 2)$; and so on.

Throughout the paper, we use a simple color scheme to represent the agent's behavior. If a node is white, this means that, at that node, the agent agrees to play a 50:50 bet. If the node is black, this means that the agent does *not* play a 50:50 bet at that node, either because he leaves the casino when he arrives at that node, or because his actions in earlier rounds prevent him from even reaching that node. For example, the interpretation of Figure 2 is that the agent agrees to enter the casino at time 0 and then keeps gambling until time $T = 5$ or until he hits node $(3, 1)$, whichever comes first. The fact that node $(3, 1)$ has a black color immediately implies that node $(4, 1)$ must also have a black color: a node that can only be reached by passing through a black node must itself be black.

As noted above, the basic gamble offered by the casino in our model is a 50:50 bet to

[9] We discuss an infinite horizon analog of the finite horizon model in Section 4.4.

win or lose \$$h$. We assume that the gain and the loss are equally likely only because this simplifies the exposition, not because it is necessary for our analysis. Indeed, we have studied the case where, as in actual casinos, the basic gamble has a somewhat negative expected value – for example, where it entails a 0.46 chance of winning \$$h$, say, and a 0.54 chance of losing \$$h$ – and find that the results are similar to those that we present below.

Now that we have described the structure of the casino, we are ready to present the behavioral assumption that drives our analysis. Specifically, we assume that, at each moment of time, the agent in our model decides what to do by *maximizing the cumulative prospect theory value of his accumulated winnings or losses at the moment he leaves the casino*, where the cumulative prospect theory value of a distribution is given by (3)-(6).

In any application of prospect theory, a key step is to specify the argument of the prospect theory value function $v(\cdot)$, in other words, the "gain" or "loss" that the agent applies the value function to. As noted in the previous paragraph, our assumption is that, at each moment of time, the agent applies the value function to his *overall* winnings at the moment he leaves the casino. In the language of "reference points," our assumption is that, throughout the evening of gambling, the agent's reference point remains fixed at his initial wealth when he entered the casino, so that the argument of the value function is his wealth when he leaves the casino minus his wealth when he entered.

Our modeling choice is motivated by the way people discuss their casino experiences. If a friend or colleague tells us that he recently went to a casino, we tend to ask him "How much did you win?," not "How much did you win last year in all your casino visits?" or "How much did you win in each of the games you played at the casino?" In other words, it is overall winnings during a single casino visit that seem to be the focus of attention.

Our behavioral assumption immediately raises an important issue, one that plays a central role in our analysis. This is the fact that cumulative prospect theory – in particular, its probability weighting feature – generates a time inconsistency: the agent's *plan*, at time t, as to what he would do if he reached some later node is not necessarily what he actually does when he reaches that node.

To see the intuition, consider the node indicated by an arrow in the upper part of the tree in Figure 2, namely node $(4,1)$ – ignore the specific black or white node colorations – and suppose that the per-period bet size is $h = \$10$. It is possible to check that, *from the perspective of time 0*, the agent's preferred plan, for almost all preference parameter values, is to gamble in node $(4,1)$, should he arrive in that node. The reason is that, by gambling in node $(4,1)$, he gives himself a chance of leaving the casino in node $(5,1)$ with an overall gain of \$50. From the perspective of time 0, this gain has low probability, namely $\frac{1}{32}$, but, under cumulative prospect theory, low tail probabilities are overweighted, making node $(5, 1)$ very appealing to the agent. In spite of the concavity of the value function $v(\cdot)$ in the region of gains, then, his preferred plan, as of time 0, is almost always to gamble in node (4,1), should he reach that node.

While the agent's preferred plan, as of time 0, is to gamble in node $(4, 1)$, it is easy to see that, if he actually arrives in node $(4, 1)$, he will instead stop gambling, contrary to his initial plan. If he stops gambling in node $(4, 1)$, he leaves the casino with an overall gain of \$40. If he continues gambling, he has a 0.5 chance of an overall gain of \$50 and a 0.5 chance of an overall gain of \$30. He therefore leaves the casino in node $(4, 1)$ if

$$v(40) \geq v(50)w(\frac{1}{2}) + v(30)(1 - w(\frac{1}{2})); \tag{7}$$

in words, if the cumulative prospect theory value of leaving is greater than or equal to the cumulative prospect theory value of staying. Condition (7) simplifies to

$$v(40) - v(30) \geq (v(50) - v(30))w(\frac{1}{2}). \tag{8}$$

It is straightforward to check that condition (8) holds for *all* α, $\delta \in (0, 1)$, so that the agent indeed leaves the casino in node $(4, 1)$, contrary to his initial plan. What is the intuition? From the perspective of time 0, node $(5, 1)$ was unlikely, overweighted, and hence appealing. From the time 4 perspective, however, it is no longer unlikely: once the agent is at node $(4, 1)$, the probability of reaching node $(5, 1)$ is 0.5. The weighting function $w(\cdot)$ *underweights* moderate probabilities like 0.5. This, together with the concavity of $v(\cdot)$ in the region of gains, means that, from the perspective of time 4, node (5,1) is no longer as appealing. The agent therefore leaves the casino in node $(4, 1)$.

There is an analogous and, as we will see later, more important time inconsistency in the bottom part of the tree. For example, for almost all preference parameter values, the agent's preferred plan, from the perspective of time 0, is to stop gambling in node $(4, 5)$ – the node indicated by an arrow in the bottom part of the tree in Figure 2 – should he arrive in this node. However, if he actually arrives in node $(4, 5)$, he keeps gambling, contrary to his initial plan. The intuition for this inconsistency parallels the intuition for the inconsistency in the upper part of the tree.

Given the time inconsistency, the agent's behavior depends on two things. First, it depends on whether he is aware of the inconsistency. An agent who *is* aware of the inconsistency has an incentive to try to commit to his initial plan of action. For this agent, then, his behavior further depends on whether he is indeed able to commit. To explore these distinctions, we consider three types of agents. Our classification parallels the one used in the literature on hyperbolic discounting.

The first type of agent is "naive". An agent of this type is not aware of the time inconsistency generated by probability weighting. We analyze his behavior in Section 3.1. The second type of agent is "sophisticated" – he is aware of the time inconsistency – but is unable to find a way of committing to his initial plan. We analyze his behavior in Section 3.2. The third and final type of agent is also sophisticated – he is also aware of the time inconsistency – but is able to find a way of committing to his initial plan. We analyze his behavior in Section 3.3.[10]

3.1 Case I: The naive agent

We analyze the naive agent's behavior in two steps. First, we study his time 0 decision as to whether to enter the casino. If we find that, for some preference parameter values, he is willing to enter, we then look, for these parameter values, at his behavior *after* he enters, in other words, at his behavior for $t > 0$.

The initial decision

At time 0, the naive agent chooses a plan of action. A "plan" is a mapping from each node in the binomial tree between $t = 0$ and $t = T - 1$ to one of two possible actions: "exit," which indicates that the agent plans to leave the casino if he arrives at that node; or "continue," which indicates that he plans to keep gambling if he arrives at that node. We denote the set of all possible plans as $S_{(0,1)}$, where the subscript indicates that this is the set of plans that is available in node $(0, 1)$, the left-most node in the tree. The set $S_{(0,1)}$ grows rapidly in size as T increases: even for $T = 5$, the number of possible plans is large.[11]

For each plan $s \in S_{(0,1)}$, there is a random variable $\widetilde{G}_s$ that represents the accumulated winnings or losses the agent will experience if he exits the casino at the nodes specified by

[10] In his classic discussion of non-expected utility preferences, Machina (1989) identifies three kinds of agents: β-types, γ-types, and δ-types. In the context of our casino, the behavior of these three types of agents would be identical to the behavior of the naive agents, the sophisticates who are able to commit, and the sophisticates who are unable to commit, respectively.

[11] Since, for each of the $T(T + 1)/2$ nodes between time 0 and time $T - 1$, the agent can either exit or continue, the number of plans in $S_{(0,1)}$ is, in principle, equal to 2 to the power of $T(T + 1)/2$. The number of *distinct* plans is much lower, however. For example, for any $T \geq 3$, all plans that assign the action "exit" to node $(0, 1)$ are effectively the same, as are all plans that assign the actions "continue," "exit," and "exit" to nodes $(0, 1)$, $(1, 1)$, and $(1, 2)$, respectively.

plan s. For example, if s is the exit strategy shown in Figure 2, then

$$\widetilde{G}_s \sim (\$30, \frac{7}{32}; \$10, \frac{9}{32}; -\$10, \frac{10}{32}; -\$30, \frac{5}{32}; -\$50, \frac{1}{32}).$$

With this notation in hand, we can write down the problem that the naive agent solves at time 0. It is:

$$\max_{s \in S_{(0,1)}} V(\widetilde{G}_s), \tag{9}$$

where $V(\cdot)$ computes the cumulative prospect theory value of the gamble that is its argument. Suppose that $V(\widetilde{G}_s)$ attains its maximum value for plan $s^* \in S_{(0,1)}$. The naive agent then enters the casino – in other words, he plays a gamble at time 0 – if and only if $V^* \equiv V(\widetilde{G}_{s^*}) > 0$.[12] We emphasize that the naive agent chooses a plan at time 0 without regard for the possibility that he might stray from the plan in future periods. After all, he is naive: he does not realize that he might later depart from the plan.[13]

The nonlinear probability weighting embedded in $V(\cdot)$ makes it very difficult to solve problem (9) analytically; indeed, the problem has no known analytical solution for general T. We therefore solve it numerically, focusing on the case of $T = 5$. We are careful to check the robustness of our conclusions by solving (9) for a wide range of preference parameter values.[14]

The time inconsistency generated by probability weighting means that we cannot use backward induction to solve problem (9). Instead, we use the following procedure. For each plan $s \in S_{(0,1)}$ in turn, we compute the gamble $\widetilde{G}_s$ and calculate its cumulative prospect theory value $V(\widetilde{G}_s)$. We then look for the plan s^* that maximizes $V(\widetilde{G}_s)$ and check whether $V^* > 0$.

We begin our analysis by identifying the range of preference parameter values for which the naive agent enters the casino. We set $T = 5$, $h = \$10$, and restrict our attention to preference parameter triples $(\alpha, \delta, \lambda)$ for which $\alpha \in [0, 1]$, $\delta \in [0.3, 1]$, and $\lambda \in [1, 4]$.[15]

[12]Since $S_{(0,1)}$ includes the strategy of not entering the casino at all – this is the strategy that assigns the action "exit" to node $(0, 1)$ – the value of V^* must be at least zero, the cumulative prospect theory value of not entering. The agent enters the casino if there is a plan that involves gambling in node $(0, 1)$ whose cumulative prospect theory value is strictly greater than zero.

[13]We only allow the agent to consider path-independent plans of action: his planned action at time t depends only on his accumulated winnings at that time and not on the path by which he accumulated those winnings.

[14]For one special case, the case of $T = 2$, a full analytical characterization of the behavior of all three types of agents *is* available. We discuss this case in detail in Sections 4.3 and 7.3.

[15]The behavior of the three types of agents that we consider does not depend on the value of h; we set $h = \$10$ only for the sake of concreteness. We have also studied the case of $T = 10$ and find that the results in this case parallel those for $T = 5$. We do not use $T = 10$ as our benchmark case, however, because of its

We focus on values of λ that are less than 4 so as not to stray too far from Tversky and Kahneman's (1992) estimate of this parameter; and, as noted earlier, we restrict attention to values of δ that exceed 0.3 so as to ensure that the weighting function (6) is monotonically increasing. We then discretize each of the intervals $[0,1]$, $[0.3,1]$, and $[1,4]$ into a set of 20 equally-spaced points and study parameter triples (α,δ,λ) where each parameter takes a value that corresponds to one of the discrete points. In other words, we study the $20^3 = 8,000$ parameter triples in the set Δ, where

$$\begin{aligned}\Delta &= \{(\alpha,\delta,\lambda) : \alpha \in \{0, 0.053, \ldots, 0.947, 1\},\\ &\quad \delta \in \{0.3, 0.337, \ldots, 0.963, 1\}, \lambda \in \{1, 1.16, \ldots, 3.84, 4\}\}. \end{aligned} \tag{10}$$

The "+" and "*" signs in Figure 3 mark the preference parameter triples for which the naive agent enters the casino, in other words, the triples for which $V^* > 0$. We explain the significance of each of the two signs below – for now, the reader can ignore the distinction. To make the marked region easier to visualize, we use a color scheme in which different colors correspond to different vertical elevations. Specifically, the blue, red, green, cyan, magenta, and yellow colors correspond to parameter triples for which λ – the parameter on the vertical axis – takes a value in the intervals [1,1.5), [1.5,2), [2,2.5), [2.5,3), [3,3.5), and [3.5,4], respectively.[16] Finally, the small circle marks Tversky and Kahneman's (1992) median estimates of the preference parameters, namely

$$(\alpha,\delta,\lambda) = (0.88, 0.65, 2.25). \tag{11}$$

Figure 3 illustrates our first main result: that, even though the agent is loss averse and even though the casino offers only 50:50 bets with zero expected value, there is still a wide range of preference parameter values for which the agent *is* willing to enter the casino. In particular, he is willing to enter for 1,813 of the 8,000 parameter triples in the set Δ. Note that, for the median estimates in (11), the agent does not enter the casino. Nonetheless, for parameter values that are not far from those in (11), he is willing to enter.

To understand why, for many parameter values, the agent is willing to enter, we study his optimal exit plan s^*. Consider the case of $(\alpha,\delta,\lambda) = (0.95, 0.5, 1.5)$, a parameter triple for which the agent enters. The left panel in Figure 4 shows the optimal exit plan in this case. Recall that, if a node has a white color, the agent plans to gamble at that node, should

much greater computational demands.

[16] In topographical terms, the region marked by the "+" and "*" signs in Figure 3 consists of two "hills" – a steep hill in the right part of the figure and a gentler hill in the left part – with a valley in the center.

he reach it. By contrast, if a node has a black color, he plans not to gamble at that node. The figure shows that, roughly speaking, the agent's optimal plan is to keep gambling until time T or until he starts accumulating losses, whichever comes first.

The exit plan in the left panel of Figure 4 helps us understand why it is that, even though the agent is loss averse and even though the casino offers only zero expected value bets, the agent is still willing to enter. The reason is that, through his choice of exit plan, the agent is able to give his perceived *overall* casino experience a positively skewed distribution: by exiting once he starts accumulating losses, he limits his downside; and by continuing to gamble when he is winning, he retains substantial upside. Since the agent overweights the tails of probability distributions, he may like the positively skewed distribution offered by the overall casino experience. In particular, under probability weighting, the chance, albeit small, of leaving the casino in the top-right node $(T, 1)$ with a large accumulated gain of $\$Th$ is very enticing. In summary, then, while the agent would always turn down the basic 50:50 bet offered by the casino if that bet were offered in isolation, he is nonetheless able, through a specific choice of exit strategy, to give his perceived overall casino experience a positively skewed distribution, one which, with sufficient probability weighting, he finds attractive.[17]

The left panel in Figure 4 shows the naive agent's optimal plan when $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$. What does the optimal plan look like for the other preference parameter triples for which he enters the casino? To answer this, we introduce some terminology. We label a plan a "gain-exit" plan if, under the plan, the agent's expected length of time in the casino conditional on exiting with a gain is less than his expected length of time in the casino conditional on exiting with a loss. Put simply, a gain-exit plan is one in which the agent plans to leave quickly if he is winning but to stay longer if he is losing. Similarly, a plan is a "loss-exit" ("neutral-exit") plan if, under the plan, the agent's expected length of time in the casino conditional on exiting with a gain is greater than (the same as) his expected length of time in the casino conditional on exiting with a loss. For example, the plan in the left panel of Figure 4 is a loss-exit plan because, conditional on exiting with a loss, the agent spends only one period in the casino, while conditional on exiting with a gain, he spends five periods in the casino.

[17] A number of authors – see, for example, Benartzi and Thaler (1995) – have noted that prospect theory may be able to explain why someone would turn down a single play of a positive expected value bet – a 50:50 bet to win \$200 or lose \$100, say – but would agree to multiple plays of the bet, a pattern of behavior that is sometimes observed in practice. This is a different point from the one we are making in this paper. After all, a prospect theory agent would turn down T plays of the basic bet offered by our casino for *any* $T \geq 1$. The reason the agent enters the casino hinges on the fact that a casino with T potential rounds of gambling is *not* the same as T plays of the casino's basic bet: in the casino, the agent has the option to leave after each round of gambling.

The "*" signs in Figure 3 mark the preference parameter triples for which the naive agent enters the casino with a *loss*-exit plan in mind.[18] In particular, for 1,021 of the 1,813 parameter triples for which the naive agent enters, he does so with a loss-exit plan in mind, one that is either identical to that in the left panel of Figure 4 or else one that differs from it in only a very small number of nodes.

Figure 3 shows that the naive agent is more likely to enter the casino with a loss-exit plan for low values of δ, for low values of λ, and for high values of α. The intuition is straightforward. By adopting a loss-exit plan, the agent gives his perceived overall casino experience a positively skewed distribution. As δ falls, the agent overweights the tails of probability distributions more heavily. He is therefore more likely to find the positively skewed distribution generated by the loss-exit plan appealing. As λ falls, the agent becomes less loss averse. He is therefore less scared by the losses he could incur under a loss-exit plan and therefore more willing to enter with such a plan. Finally, as α increases, the marginal utility of additional gains diminishes less rapidly. The agent is therefore more excited about the possibility of a large win inherent in a loss-exit plan and hence more likely to enter the casino with a plan of this kind.

For 1,021 of the 1,813 parameter triples for which the naive agent enters the casino, then, he does so with a loss-exit plan in mind. For the remaining 792 parameter triples for which he enters, he does so with a *gain*-exit plan in mind, one where he plans to gamble for longer in the region of losses than in the region of gains. These parameter triples are indicated by the "+" signs in Figure 3. As the figure shows, these parameter triples lie quite far from the median estimates in (11): most of them correspond to values of α and λ that are much lower than the median estimates or to values of δ that are much higher.

Why does the naive agent sometimes enter the casino with a gain-exit plan? Under a plan of this kind, the agent's perceived casino experience has a *negatively* skewed distribution, one with a moderate probability of a small gain and a low probability of a large loss. If α is very low, the large loss is only slightly more frightening than a small loss; and if δ is very high, the low probability of the large loss is barely overweighted. As a result, when α is low and δ is high, the agent may find the negatively skewed distribution appealing.

Our analysis shows that the component of prospect theory most responsible for the agent entering the casino is the probability weighting function: for the majority of the preference parameter values for which he enters, the agent chooses a plan that corresponds to a positively skewed casino experience; and this, in turn, is attractive precisely because of the weighting

[18]In case it is not clear, the "*" signs are located in the right half of the marked region in the figure. By contrast, the left half of the marked region is made up of "+" signs.

function. Nonetheless, the naive agent may enter the casino even in the absence of probability weighting, in other words, even if $\delta = 1$. For example, he enters when $(\alpha, \delta, \lambda) = (0.5, 1, 1.2)$. For this parameter triple, the agent's optimal plan is a gain-exit plan, one that gives his perceived casino experience a negatively skewed distribution – but since $\delta = 1$ and α is so low, the agent finds it appealing.

We noted earlier that, of all casino games, it is blackjack that most closely matches the game in our model. However, our model may also explain why another casino game, the slot machine, is as popular as it is. In our framework, an agent who enters the casino usually does so because he relishes the positively skewed distribution he perceives it to offer. Since slot machines already offer a skewed payoff, they may make it easier for the agent to give his overall casino experience a significant amount of positive skewness. It may therefore make sense that they would outstrip blackjack in popularity.

Figure 3 shows the range of preference parameter values for which the naive agent enters the casino when $T = 5$. The range of preference parameter values for which he would enter a casino with $T > 5$ rounds of gambling is at least as large as the range marked in Figure 3. To see why, note that any plan that can be implemented in a casino with $T = \tau$ rounds of gambling can also be implemented in a casino with $T = \tau + 1$ rounds of gambling. If an agent is willing to enter a casino with $T = \tau$ rounds of gambling, then, he will also be willing to enter a casino with $T = \tau + 1$ rounds of gambling: at the very least, when $T = \tau + 1$, he can just adopt the plan that leads him to enter when $T = \tau$.

Can we say more about what happens for higher values of T? For example, Figure 3 shows that, when $T = 5$, the agent does not enter the casino for Tversky and Kahneman's (1992) median estimates of the preference parameters. A natural question is then: are there *any* values of T for which a naive agent with the median preference parameter values would be willing to enter the casino? The following proposition, which provides a sufficient condition for the naive agent to be willing to enter, allows us to answer this question. The proof of the proposition is in the Appendix.

Proposition 1: A naive agent with cumulative prospect theory preferences and the preference parameters $(\alpha, \delta, \lambda)$ is willing to enter a casino offering $T \geq 2$ rounds of gambling and a basic bet of $(\$h, 0.5; -\$h, 0.5)$ if[19]

$$\sum_{j=1}^{T-\left[\frac{T}{2}\right]} (T + 2 - 2j)^{\alpha} \left(w(2^{-T} \binom{T-1}{j-1}) - w(2^{-T} \binom{T-1}{j-2}) \right) > \lambda w(\frac{1}{2}). \tag{12}$$

[19] In this expression, $\binom{T-1}{-1}$ is assumed to be equal to 0.

To derive condition (12), we take one particular exit strategy which, from our numerical analysis, we know to be either optimal or close to optimal for a wide range of preference parameter values – roughly speaking, this is a strategy where the agent keeps gambling if he is winning but stops gambling if he starts accumulating losses – and compute its cumulative prospect theory value explicitly. Condition (12) checks whether this value is positive; if it is, we know that the naive agent enters the casino. While the condition is hard to interpret, it is useful because it allows us to learn something about the agent's behavior when T is high without solving problem (9), something which, for high values of T, is computationally very taxing.[20]

It is easy to check that, for Tversky and Kahneman's (1992) estimates, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$, the lowest value of T for which condition (12) holds is $T = 26$. We can therefore state the following corollary.

Corollary: If $T \geq 26$, a naive agent with cumulative prospect theory preferences and the parameter values $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$ is willing to enter a casino with T rounds of gambling and a basic bet of $(\$h, 0.5; -\$h, 0.5)$.

We noted earlier that we are dividing our analysis of the naive agent's behavior into two parts. We have just completed the first part: the analysis of the agent's time 0 decision as to whether to enter the casino. We now turn to the second part: the analysis of what the agent does at time $t > 0$.

Subsequent behavior

Suppose that, at time 0, the naive agent decides to enter the casino. In node (t, j) at some later time $t \geq 1$, he solves

$$\max_{s \in S_{(t,j)}} V(\widetilde{G}_s). \tag{13}$$

Here, $S_{(t,j)}$ is the set of plans the agent could follow from time t onward, where, in a similar way to before, a "plan" is a mapping from each node between time t and time $T - 1$ to one of two actions: "exit," indicating that the agent plans to leave the casino if he reaches that node, or "continue," indicating that the agent plans to keep gambling if he reaches that node. As before, $\widetilde{G}_s$ is a random variable that represents the agent's potential accumulated winnings or losses if he follows plan s, and $V(\widetilde{G}_s)$ is its cumulative prospect theory value.

[20] Although condition (12) is only a sufficient condition, there is a sense in which it is an accurate sufficient condition: at least for the low values of T where it is possible to check, the set of triples $(\alpha, \delta, \lambda)$ that satisfy condition (12) is very similar to the set of triples for which the naive agent actually enters the casino with a loss-exit plan. If we denote the left-hand side of condition (12) as X, it is possible to show that $w(\frac{1}{2}) > \lambda X$ is also a sufficient condition for entry. For low T, this last condition accurately approximates the set of triples for which the naive agent enters the casino with a *gain*-exit plan.

For example, if the agent is in node $(3,1)$, the plan under which he leaves at time $T = 5$, but not before, corresponds to

$$\widetilde{G}_s \sim (\$50, \frac{1}{4}; \$30, \frac{1}{2}; \$10, \frac{1}{4}).$$

If $s^* \in S_{(t,j)}$ is the plan that solves problem (13), the agent gambles in node (t,j) if and only if

$$V(\widetilde{G}_{s^*}) > v(h(t+2-2j)), \tag{14}$$

where the right-hand side is the utility of leaving the casino at this node.[21]

To see how the naive agent behaves for $t \geq 1$, we first return to the example from earlier in this section in which $T = 5$ and $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$. For these parameter values, the naive agent enters the casino at time 0. The right panel of Figure 4 shows what he does subsequently, at time $t \geq 1$. Recall that the left panel in the figure shows the initial plan of action he constructs at time 0.

Figure 4 shows that, while the naive agent's initial plan was to gamble as long as possible when winning but to stop if he started accumulating losses, he actually, roughly speaking, does the opposite: he gambles as long as possible when losing and stops once he accumulates some gains. Our model therefore captures a common intuition, namely that people often gamble more than they planned to in the region of losses.

As noted earlier, the time inconsistency is entirely driven by the probability weighting function. As of time 0, a strategy under which the agent continues to gamble if he is winning is very attractive: under this strategy, the agent could take home \$50 in node $(5,1)$. While this is unlikely, the low probability of it happening is overweighted, making node $(5,1)$ very appealing. If the agent wins the first few bets, however, reaching node $(5,1)$ is no longer an unlikely outcome, and, as such, is no longer overweighted. This, together with the concavity of the value function in the region of gains, means that the agent stops gambling after earning some gains, contrary to his initial plan.

A similar mechanism is at work in the lower part of the tree. As of time 0, a plan under which the agent continues to gamble if he is losing is very unattractive: such a plan exposes the agent to low probability losses, which, given that the agent overweights tail probabilities, is very unappealing. If the agent starts losing, however, losses that were initially unlikely

[21] Since $S_{(t,j)}$ includes the strategy of leaving the casino in node (t,j) – this is the strategy that assigns the action "exit" to node (t,j) – the value of $V(\widetilde{G}_{s^*})$ must be at least $v(h(t+2-2j))$. The agent gambles in node (t,j) if there is a plan that involves gambling in node (t,j) whose cumulative prospect theory value is strictly greater than $v(h(t+2-2j))$.

and hence overweighted are no longer unlikely and therefore no longer overweighted. This, together, with the convexity of the value function in the region of losses, means that the agent continues to gamble if he is losing, contrary to his initial plan.[22]

How typical is the behavior in the right panel of Figure 4? Earlier in this section, we described a numerical analysis of 8,000 preference parameter triples and noted that the naive agent enters the casino for 1,813 of these 8,000 triples. We find that, for *all* 1,813 of these triples, the agent's actual behavior in the casino is described by a gain-exit strategy that is either exactly equal to the one in the right panel of Figure 4 or else one that differs from it in only a very small number of nodes; indeed, for triples for which $\alpha > 0$, the agent *always* gambles until $T = 5$ in the region of losses. We noted earlier that, for 1,021 of the 1,813 triples for which the naive agent enters the casino, his initial plan is a loss-exit plan. In all 1,021 of these cases, then, the naive agent's actual behavior is, roughly speaking, the opposite of what he initially planned.[23]

3.2 Case II: The sophisticated agent, without commitment

In section 3.1, we considered the case of a naive agent – an agent who is unaware of the time inconsistency generated by probability weighting. In Sections 3.2 and 3.3, we study sophisticated agents, in other words, agents who are aware of the time inconsistency. A sophisticated agent has an incentive to try to commit to his time 0 plan. In this section, we consider the case of a sophisticated agent who is *unable* to find a way of committing to his time 0 plan; we label this agent a "no-commitment sophisticate". In Section 3.3, we study the case of a sophisticated agent who *is* able to commit to his initial plan.

To decide on a course of action, the no-commitment sophisticate uses backward induction, working leftward from the right-most column of the binomial tree. If he has not yet left the casino at time T, he must exit at that time. Knowing this, he is able to determine what he will do at time $T-1$. This, in turn, allows him to determine what he will do at time $T-2$, and so on.

[22]The naive agent's "naivete" can be interpreted in two ways. The agent may fail to realize that, after he starts gambling, he will be tempted to depart from his initial plan. Alternatively, he may recognize that he will be tempted to depart from his initial plan, but he may erroneously think that he will be able to resist the temptation. After many casino visits, the agent may learn his way out of the first kind of naivete. It may take him much longer, however, to learn his way out of the second kind. People often continue to believe that they will be able to exert self-control in the future even when they have repeatedly failed to do so in the past.

[23]For the remaining 792 triples for which the naive agent enters the casino, his actual behavior is more similar to his initial plan: both his initial plan and actual behavior are gain-exit strategies in which he gambles longer in the region of losses than in the region of gains.

Mathematically, the no-commitment sophisticate gambles in node (t, j), where $t \in [0, T-1]$, if and only if

$$V(\widetilde{G}_{t,j}) > v(h(t + 2 - 2j)). \qquad (15)$$

The term $v(h(t+2-2j))$ is the utility of leaving the casino in node (t, j). The term $V(\widetilde{G}_{t,j})$ is the value of continuing to gamble: specifically, it is the cumulative prospect theory value of the random variable $\widetilde{G}_{t,j}$ which represents the accumulated winnings or losses the agent will exit the casino with if he gambles in node (t, j). The random variable $\widetilde{G}_{t,j}$ is determined by the exit strategy computed in earlier steps of the backward iteration. Note that, precisely because he computes his course of action using backward induction, the no-commitment sophisticate is time consistent.

We study the behavior of the no-commitment sophisticate for $T = 5$ and for each of the 8,000 preference parameter triples in the set Δ defined in (10). The "+" signs in Figure 5 mark the triples for which the no-commitment sophisticate enters the casino – in other words, the triples for which $V(\widetilde{G}_{0,1}) > 0$. As before, we use a color scheme to make the marked region easier to visualize. The blue, red, green, and cyan colors correspond to parameter triples for which λ lies in the intervals [1,1.5), [1.5,2), [2,2.5), and [2.5,4], respectively.

The figure shows that the agent enters the casino for only a narrow range of parameter triples: specifically, for just 753 parameter triples, all of which lie far from the small circle which marks Tversky and Kahneman's (1992) estimates in (11). The intuition is straightforward. The agent realizes that, if he does enter the casino, he will gamble longer in the region of losses than in the region of gains. This will give his overall casino experience a *negatively* skewed distribution. Since he overweights the tails of distributions, he usually finds this unattractive and therefore refuses to enter. Reasoning of this type may in part explain why the majority of Americans do *not* gamble in casinos.

For the 753 parameter triples for which the no-commitment sophisticate enters the casino, he follows a gain-exit strategy. When α and λ are sufficiently low and δ is sufficiently high, the negatively skewed casino experience generated by this strategy is actually appealing.

3.3 Case III: The sophisticated agent, with commitment

In this section, we study the behavior of a sophisticated agent who is able to commit to his initial plan. We call this agent a "commitment-aided sophisticate."

We proceed in the following way. We assume that, at time 0, the agent can find a way of committing to *any* exit strategy $s \in S_{(0,1)}$. Once we identify the strategy that he would choose, we then discuss how he might actually commit to this strategy in practice.

At time 0, then, the commitment-aided sophisticate solves:

$$\max_{s \in S_{(0,1)}} V(\widetilde{G}_s). \tag{16}$$

In particular, since he can commit to any exit strategy, we do not need to restrict the set of strategies he considers. He searches across *all* elements of $S_{(0,1)}$ until he finds the strategy s^* with the highest cumulative prospect theory value $V^* = V(\widetilde{G}_{s^*})$. He enters the casino if and only if $V^* > 0$.

The problem in (16) is identical to the problem solved by the naive agent at time 0. The two types of agents therefore enter the casino for exactly the same range of preference parameter values. For $T = 5$, for example, the commitment-aided sophisticate enters the casino for the 1,813 parameter triples marked by the "+" and "*" signs in Figure 3. Moreover, for any given parameter triple, the commitment-aided sophisticate and the naive agent enter the casino with exactly the same strategy in mind. For example, for the 1,021 parameter triples indicated by "*" signs in Figure 3, the commitment-aided sophisticate enters the casino with a loss-exit plan in mind, as does the naive agent.

The naive agent and the commitment-aided sophisticate solve the same problem at time 0 because they both *think* that they will be able to maintain any plan they select at that time. The two types of agents differ, however, in what they do after they enter the casino. Since he has a commitment device at his disposal, the commitment-aided sophisticate is able to stick to his initial plan. The naive agent, on the other hand, deviates from his initial plan.

For the 1,021 parameter triples indicated by "*" signs in Figure 3, then, the commitment-aided sophisticate would like to commit to a loss-exit strategy. The question now is: how does he commit to such a strategy? For example, in the lower part of the binomial tree, how does he manage to stop gambling when he is losing even though he is tempted to continue? And in the upper part of the tree, how does he manage to continue gambling when he is winning even though he is tempted to stop?[24]

In the lower part of the tree, one simple commitment strategy is for the agent to go to the casino with only a small amount of cash and to leave his ATM card at home. If he starts accumulating losses, he is sorely tempted to continue gambling, but, since he has run out of cash, he has no option but to go home. It is a prediction of our model that some gamblers

[24] We are focusing on the parameter triples marked by "*" signs in Figure 3 because they are closer to Tversky and Kahneman's (1992) median parameter estimates than are the triples marked by "+" signs. Note also that, for the latter set of triples, the commitment problem is less severe because the agent's initial preferences are more aligned with his subsequent preferences: both preferences favor gain-exit strategies.

will use a strategy of this kind. Anecdotal reports suggest that this *is* a common gambling strategy.

It is harder to think of a common strategy that gamblers use to solve the commitment problem in the upper part of the tree. In a way, this is not surprising. An interesting prediction of our model – a prediction that we have found to hold all the more strongly for higher values of T – is that the time inconsistency is more severe in the lower part of the tree than in the upper part. Comparing the two panels in Figure 4, we see that, in the lower part of the tree, the time inconsistency, and hence the commitment problem, is severe: the agent wants to gamble in *every* node in the region of losses even though his initial plan was to gamble in none of them. In the upper part of the tree, however, the time inconsistency, and hence the commitment problem, is less acute: the agent's initial plan conflicts with his subsequent preferences in only a few nodes. It therefore makes sense that the commitment strategies gamblers use in practice seem to be aimed primarily at the time inconsistency in the lower part of the tree.

Although it is hard to think of ways in which gamblers commit to their initial plan in the upper part of the tree, note that here, casinos have an incentive to help. In general, casinos offer bets with negative expected values; it is therefore in their interest that gamblers stay on site as long as possible. From the casinos' perspective, it is alarming that gamblers are tempted to leave earlier than planned when they are winning. This may explain the practice among some casinos of offering vouchers for free food and lodging to people who are winning. In our framework, casinos do this in order to encourage gamblers who are thinking of leaving with their gains, to stay longer.[25]

4 Further Remarks

4.1 Average losses

Our analysis shows that the set of casino gamblers consists primarily of two distinct types: naive agents and commitment-aided sophisticates. Which of these two types loses more

[25]When the commitment-aided sophisticate chooses a plan to commit to at time 0, he does not put any weight on the preferences of his future "selves." For example, when $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$, he commits to the plan in the left panel of Figure 4 even though he knows that, in node $(1, 2)$, his time 1 self will want to continue gambling. One interpretation of this is that the time 0 self disapproves of his future preferences. This, in turn, can be justified by arguing that, in wanting to deviate from the time 0 plan, the time 1 self is committing an error of "consequentialism": of ignoring risk that was previously borne but not realized (Machina, 1989). Nonetheless, it may be interesting to study a model in which the time 0 self puts at least some weight on future selves' preferences when choosing a plan to commit to.

money in the casino, on average?

In the context of the model of Section 3, the answer is straightforward. Since the basic bet offered by the casino has an expected value of zero, average winnings are zero for both types. However, if, as in actual casinos, the basic bet has a negative expected value, then an agent's average winnings are the (negative) expected value of the basic bet multiplied by the average number of rounds the agent gambles. To determine whether the naive agent loses more, on average, than the commitment-aided sophisticate, we therefore need to check whether he gambles for longer, on average.

We have studied the gambling behavior of the two types of agents when the basic bet offered by the casino has a negative expected value – for example, when it takes the form $(\$h, 0.46; -\$h, 0.54)$. We find that, for most preference parameter values – although not all – the naive agent stays in the casino longer, on average, than does the sophisticated agent. In most cases, then, his average losses are larger.

4.2 Competition from lotteries

A natural question raised by our analysis is: how can casinos survive competition from lottery providers? After all, the one-shot gambles offered by lottery providers may be a more convenient source of the positive skewness that the casino goers in our framework are seeking.

In this section, we discuss one mechanism through which casinos can survive competition from lotteries – a mechanism that we can analyze using the framework of Section 3. We demonstrate the idea formally with the help of a simple equilibrium model, presented in detail in Section 7.2 in the Appendix.

In this model, there is competitive provision of both one-shot lotteries and casinos and both types of firms incur a cost per consumer served. Even so, both lottery providers *and* casinos manage to break even. In equilibrium, lottery providers attract the no-commitment sophisticates. These agents prefer lotteries to casinos because they know that, in a casino, they would face a negatively skewed, and hence unattractive, distribution of accumulated gains and losses.

Casinos compete with lottery providers by offering slightly better, albeit unfair, odds. This attracts the commitment-aided sophisticates and the naive agents, both of whom think that, through a particular choice of exit strategy, they can construct an overall casino experience – in other words, a distribution of accumulated gains and losses – whose prospect theory value exceeds the prospect theory value offered by one-shot lotteries. The commitment-aided sophisticates are indeed able to construct such a casino experience. Since casinos incur a cost per consumer served, they lose money on these agents. They make these losses up, however, on the naive agents, because, as noted in Section 4.1, these agents gamble longer, on average, than they were planning to. In short, then, casinos are able to compete with lottery providers because their dynamic structure allows them to exploit naive agents' time inconsistency.

The equilibrium model in the Appendix also sheds light on a related question, namely whether casinos would want to explicitly offer, in the form of a one-shot gamble, the overall casino experience that their customers are trying to construct dynamically. According to the model, casinos would *not* want to offer such a one-shot gamble. If they did, naive agents, believing themselves to be indifferent between the one-shot and dynamic gambles, might switch to the one-shot gamble, thereby effectively converting themselves from naive agents to commitment-aided sophisticates. Casinos would then lose money, however, because it is precisely naive agents' time inconsistency that allows them to break even.

4.3 The case of $T = 2$

The decision problems of the naive agent, the no-commitment sophisticate, and the commitment-aided sophisticate – problems (9), (15), and (16) – have no known analytical solution for general T. We have therefore solved them numerically for a wide range of preference parameter values. When $T = 2$, however, these problems *do* have analytical solutions. We present these solutions in Proposition 2 in Section 7.3 of the Appendix.

Proposition 2 shows that the case of $T = 2$ is quite rich: several of the patterns that emerged in Section 3 from our analysis of the $T = 5$ case turn out to hold even when $T = 2$. For example, in Section 3 we saw that, when $T = 5$, the naive agent either does not enter the casino at all, or else enters with a loss-exit strategy in mind but actually follows a gain-exit strategy, or else enters with a gain-exit strategy in mind and then indeed follows a gain-exit strategy. We also saw that the no-commitment sophisticate enters the casino for only a limited range of preference parameter values, a range that is far removed from Tversky and Kahneman's (1992) median estimates; and that when he enters, he follows a gain-exit strategy. Proposition 2 shows that these results hold even when $T = 2$.

The proposition also shows that, in other ways, the case of $T = 2$ is *less* rich than the case of $T > 2$. In Section 4.1, we saw that, when $T = 5$, the naive agent typically stays in the casino longer, on average, than he originally planned to, a result that played a crucial role in our discussion of how casinos compete with lotteries. Proposition 2 shows, however, that when $T = 2$, the naive agent stays in the casino exactly as long, on average, as he planned to.

A more serious limitation of the $T = 2$ case is that, when $T = 2$, the naive agent enters the casino for only a very narrow range of preference parameter values. The reason is that, when $T = 2$, it is difficult for the agent to create a casino experience that is positively skewed enough to overcome the aversion to gambling that stems from loss aversion. As a result, we cannot use the $T = 2$ case to build a strong case for prospect theory as a driver of casino gambling. It is only when we study the case of $T > 2$, as we do in Section 3, that we find that the naive agent is willing to enter the casino for a wide range of preference parameter values, a range that, for sufficiently high T, includes even Tversky and Kahneman's (1992) median estimates. This allows us to build a stronger case for prospect theory as a driver of casino gambling.

4.4 The infinite horizon case

In our analysis so far, we have imposed an exogeneous date, date T, at which the agent must leave the casino if he has not already done so. We make this assumption because we think that it captures an important feature of gambling: most people simply cannot stay in a casino indefinitely. It is nonetheless interesting to ask whether the conclusions of Section 3 also hold in an infinite horizon setting. We therefore briefly discuss this case. In short, we find that the predictions of the infinite horizon analysis are broadly consistent with those of the finite horizon analysis.

Suppose that a cumulative prospect theory agent is evaluating a casino of the kind described in Section 3, except that there is now no final date T. Without further constraints, the problem is not well-posed: the agent can achieve arbitrarily high utility by planning to exit the casino only when his accumulated gains reach $\$ah$, where a is a sufficiently large positive integer. A natural constraint to add is a limited liability constraint, one that requires the agent to leave the casino if his accumulated losses reach $\$\underline{b}h$, where $\underline{b}$ is a positive integer. Even with this constraint, however, the problem is still not well-posed for a range of empirically relevant preference parameter values. The problem becomes well-posed, however, if we change the basic bet offered by the casino from a 50:50 bet to win or lose $\$h$ to

$$(\$h, p; -\$h, 1-p)\,, \tag{17}$$

where $p < 0.5$.[26]

We therefore study an infinite horizon casino that offers the basic bet in (17), with the additional constraint that the agent must exit if his losses reach $\$\underline{b}h$. At time 0, the naive agent solves

$$\max_{s \in S_0} V(\widetilde{G}_s). \tag{18}$$

This parallels the decision problem in (9), but with some differences. S_0 is the set of strategies available to the agent when, as at time 0, his accumulated gains equal 0. We consider strategies in which the agent leaves the casino if his accumulated gains reach $\$ah$, where a is any positive integer, or if his accumulated losses reach $\$bh$, where b is any non-negative integer with $b \leq \underline{b}$. As before, $\widetilde{G}_s$ is a random variable that represents the accumulated gains or losses the agent will experience if he follows strategy $s \in S_0$, while $V(\cdot)$ computes the cumulative prospect theory value of its argument. Using standard results – see Feller (1968) – we know that

$$\widetilde{G}_s \sim (\$ah, p_a; -\$bh, p_b), \tag{19}$$

where

$$p_a = \frac{\left(\frac{1-p}{p}\right)^b - 1}{\left(\frac{1-p}{p}\right)^{a+b} - 1}, \qquad p_b = 1 - p_a. \tag{20}$$

The agent enters the casino if $V(\widetilde{G}_{s^*}) > 0$, where s^* is the plan that solves problem (18).

We solve problem (18) for $h = \$10$, $\underline{b} = 100$, $p = 0.45$, and for the same benchmark preference parameter values as before, namely $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$. We find that the naive agent enters the casino, and that the optimal gambling plan sets $a = 7$ and $b = 1$. In other words, the agent plans to gamble until he wins \$70 or loses \$10. This optimal plan has a very similar flavor to the naive agent's optimal plan in the finite horizon case, in that it generates a perceived casino experience that is positively skewed. Indeed, as before, it is this positive skewness that the agent finds attractive and that leads him to enter the casino.[27]

[26] A decision problem is not well-posed – in other words, it is ill-posed – if the supremum of the utilities that the agent can obtain from all feasible strategies is infinite. Suppose that an agent enters an infinite horizon casino offering 50:50 bets to win or lose $\$h$ with the plan to leave only when his accumulated gains reach $\$ah$ for some positive integer a. We know from standard theory that he will reach this goal with probability 1. The cumulative prospect theory value of this plan is therefore $v(ah)$, which can be arbitrarily large. With a little more algebra, we can show that, even in the presence of a limited liability constraint, the decision problem is ill-posed when $\alpha > \delta$, an empirically important case.

[27] We find that, for $p = 0.45$, the naive agent enters the casino for 1,007 of the 8,000 parameter triples in the set Δ in (10). For 309 of these 1,007 triples, he enters the casino with a loss-exit plan in mind, while for the remaining 698 triples, he enters with a gain-exit plan, where, in the context of the infinite horizon model, we define a gain-exit (loss-exit) plan as one in which the agent is more likely to leave the casino with an accumulated gain (loss) than with an accumulated loss (gain). While gain-exit plans are

What does the naive agent do subsequently? Suppose that, at some later date, his accumulated gains equal $\$kh$, for some integer k. At this point, he continues to gamble if

$$\max_{s \in S_k} V(\widetilde{G}_s) > v(kh). \tag{21}$$

This condition parallels condition (14). Here, S_k is the set of strategies available to the agent when his accumulated gains equal $\$kh$. We consider strategies for which the agent continues to gamble until his accumulated gains reach $\$ah$, where a is an integer that satisfies $a \geq k$, or until his accumulated losses reach $\$bh$, where b is an integer that satisfies $-(k-1) \leq b \leq \underline{b}$.

We find that, just as in the finite horizon case, and for similar reasons, the naive agent is time inconsistent. For example, we saw that, when $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$, the agent's initial plan was to leave the casino if he accumulated \$10 in losses. However, if he *actually* loses \$10, he instead continues to gamble, contrary to his initial plan. Indeed, we find that, in the region of losses, the naive agent only leaves the casino if forced to, in other words, only if his accumulated losses reach $\$\underline{b}h$. Just as in the finite horizon case, then, the naive agent gambles much longer than planned in the region of losses.

4.5 Predictions and other evidence

Our model makes a number of novel predictions – predictions that, we hope, will eventually be tested. Perhaps the clearest prediction is that gamblers' planned behavior will differ from their actual behavior in systematic ways. Specifically, if we survey people when they first enter a casino as to what they *plan* to do and then look at what they actually do, we should find that, on average, they exit sooner than planned in the region of gains and later than planned in the region of losses. Moreover, if gamblers who are more sophisticated in the real-world sense of the word – in terms of education or income, say – are also more sophisticated in terms of recognizing their potential time inconsistency, we should see a larger difference between planned and actual behavior among the less sophisticated.

Some experimental evidence already available in the literature gives us hope that these predictions will be confirmed in the field. Barkan and Busemeyer (1999) and Andrade and Iyer (2009) offer subjects a sequence of 50:50 bets in a laboratory setting; but before playing the gambles, subjects are asked how they *plan* to gamble in each round. Both studies find that, consistent with our model, subjects systematically gamble more than planned after an

more common overall, loss-exit plans are more common for preference parameter values that are closer to the median estimates of Tversky and Kahneman (1992).

early loss; and in Barkan and Busemeyer's study, they also gamble less than planned after an early gain.[28]

5 Conclusion

In this paper, we present a new model of casino gambling, one that is rooted in the probability weighting component of cumulative prospect theory. In recent years, probability weighting has been linked to a wide range of economic phenomena. For example, Barberis and Huang (2008) suggest that it is responsible for several empirical patterns in financial markets, including the low long-term average return on IPO stocks and the apparent overpricing of out-of-the-money options. Taken together with this prior research, then, our paper suggests that casino gambling is not an isolated phenomenon requiring its own unique explanation, but rather that it is one of a family of empirical facts, all of which are driven by the same underlying mechanism: probability weighting.

[28]If we reinterpret the binomial tree in Section 3 as representing not a casino, but rather the evolution of a stock price over time, then the results in that section also suggest predictions about how a cumulative prospect theory investor would trade a *stock* over time. We have studied this question in a framework similar to that of Section 3 and find that the model indeed delivers a range of novel predictions. For example, a naive investor, who is unaware of the time inconsistency generated by probability weighting, may exhibit a "disposition effect" in his trading – see Odean (1998) – even though he planned to exhibit the *opposite* of the disposition effect. This particular contrast between planned and actual behavior has not been noted in the prior work on the trading of prospect theory investors because this literature has ignored the dynamic effects of probability weighting.

6 References

Andrade E. and G. Iyer (2009), "Planned and Actual Betting in Sequential Gambles," *Journal of Marketing Research* 46, 372-383.

Barberis N. and M. Huang (2008), "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," *American Economic Review* 98, 2066-2100.

Barkan, R. and J.R. Busemeyer (1999), "Changing Plans: Dynamic Inconsistency and the Effect of Experience on the Reference Point," *Psychonomic Bulletin & Review* 6, 547-554.

Benartzi, S. and R. Thaler (1995), "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 75-92.

Conlisk, J. (1993), "The Utility of Gambling," *Journal of Risk and Uncertainty* 6, 255-275.

Feller W. (1968), *An Introduction to Probability Theory and Its Applications*, Wiley.

Friedman, M. and L. Savage (1948), "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy* 56, 279-304.

Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263-291.

Machina, M. (1989), "Dynamic Consistency and Non-expected Utility Models of Choice under Uncertainty," *Journal of Economic Literature* 27, 1622-1668.

Odean T. (1998), "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance* 53, 1775-1798.

Tversky, A. and D. Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5, 297-323.

7 Appendix

7.1 Proof of Proposition 1

Through extensive numerical analysis, we find that, when the naive agent enters the casino, he often chooses the following strategy or one similar to it: he exits (i) if he loses in the first round; (ii) if, after the first round, his accumulated winnings ever drop to zero; and (iii) at time T, if he has not already left by that point. Condition (12) simply checks whether the cumulative prospect theory value of this specific exit strategy is positive. If it is, we know that the agent enters the casino.

If the agent exits because he loses in the first round, then, since the payoff of $-\$h$ is the only negative payoff he can receive under the above exit strategy, its contribution to the cumulative prospect theory value of the strategy is $-\lambda h^\alpha w(0.5)$. If he exits because, at some point after the first round, his accumulated winnings drop to zero, this contributes nothing to the cumulative prospect theory value of the exit strategy, precisely because the payoff is zero. All that remains, then, is to compute the component of the cumulative prospect theory value of the exit strategy that stems from the agent exiting at date T.

Under the above exit strategy, there are $T - [\frac{T}{2}]$ date T nodes with positive payoffs at which the agent might exit, namely nodes (T, j), where $j = 1, \ldots, T - [\frac{T}{2}]$. The payoff in node (T, j) is $(T + 2 - 2j)h$. We need to compute the probability that the agent exits in node (T, j), in other words, the probability that he moves from the initial node $(0, 1)$ to node (T, j) without losing in the first round and without his accumulated winnings dropping to zero at any point after that. With the help of the reflection principle – see Feller (1968) – we compute this probability to be

$$2^{-T}\left[\binom{T-1}{j-1} - \binom{T-1}{j-2}\right].$$

The probability weight associated with node (T, j) is therefore

$$w(2^{-T}\binom{T-1}{j-1}) - w(2^{-T}\binom{T-1}{j-2}).$$

In summary then, the exit strategy we described above has positive cumulative prospect theory value – and hence the naive agent is willing to enter the casino – if

$$\sum_{j=1}^{T-[\frac{T}{2}]} ((T+2-2j)h)^\alpha \left(w(2^{-T}\binom{T-1}{j-1}) - w(2^{-T}\binom{T-1}{j-2})\right) - \lambda h^\alpha w(0.5) > 0.$$

This is condition (12).

By appropriately modifying the above argument, it is straightforward to check that, if the basic bet offered by the casino is $(\$h, p; -\$h, 1-p)$ for some $p \in (0, 1)$ rather than $(\$h, 0.5; -\$h, 0.5)$, a sufficient condition for the naive agent to enter is

$$\sum_{j=1}^{T-[\frac{T}{2}]} (T+2-2j)^\alpha \left(w(\sum_{l=1}^{j} p_l) - w(\sum_{l=1}^{j-1} p_l)\right) - \lambda w(1-p) > 0$$

where

$$p_l = p^{T-l+1}(1-p)^{l-1}\left[\binom{T-1}{l-1} - \binom{T-1}{l-2}\right].$$

7.2 A model with competitive provision of both lotteries and casinos

In this section, we show that casinos can survive in an economy with competitive provision of both lotteries and casinos. Consider an economy with two kinds of firms: "casinos" and "lottery providers." There are many firms of each kind; we index casinos with the subscript i and lottery providers with the subscript j.

Each casino has the form described in Section 3, with one exception. As before, each casino offers T rounds of gambling, but the basic bet in casino i is now $(\$h, p_i; -\$h, 1-p_i)$, where p_i is no longer necessarily equal to 0.5 but can instead take any value in the interval $(0, 0.5]$. The parameters T and $\$h$ are fixed across casinos, but each casino chooses its own value of p_i.

Lottery provider j offers consumers a one-shot gamble $\widetilde{L}_j$ of its own choosing. To keep the model tractable, we require that $\widetilde{L}_j$ satisfies the following condition: it must be possible to dynamically construct $\widetilde{L}_j$, using some exit strategy, in a hypothetical casino that offers T rounds of gambling and a basic bet of the form $(\$h, q_j; -\$h, 1-q_j)$ for some $q_j \in (0, 0.5]$.[29]

There is a continuum of consumers with a total mass of one. All consumers have the cumulative prospect theory preferences in (3)-(6) with identical preference parameters α, δ, and λ. Each consumer must either play in one of the casinos, take one of the one-shot gambles offered by lottery providers, or do nothing. He chooses the option with the highest cumulative prospect theory value. A fraction $\mu_N \geq 0$ of consumers are naive about the time inconsistency they would experience in a casino; a fraction $\mu_{S,NC} \geq 0$ are sophisticated about the time inconsistency but do not have access to a commitment device; and a fraction $\mu_{S,CA} = 1 - \mu_N - \mu_{S,NC} \geq 0$ are also sophisticated about the time inconsistency and do have access to a commitment device. Each casino and each lottery provider incurs a cost $C > 0$ per unit of consumers it serves. It is straightforward to extend our analysis to the case where

[29]The intuition of this section does not depend on the specific structure we impose on the gambles offered by lottery providers; we impose this assumption only to simplify the model. It *is* important, however, that there be a bound on the maximum loss that a lottery provider or a casino can impose on a consumer; otherwise, both lottery providers and casinos could offer consumers gambles with negative expected values but infinite cumulative prospect theory values. This is a consequence of the fact that the prospect theory value function is convex even for large losses. In a more general model that imposes risk aversion for large losses, there would be no need for an exogeneous bound on the size of a loss: consumers would simply turn down gambles with large potential losses.

casinos and lottery providers have different cost structures.

In this economy, a competitive equilibrium consists of a set $\{p_i\}$, where p_i is the win probability of the basic bet in casino i, and a set $\{\widetilde{L}_j\}$, where $\widetilde{L}_j$ is the one-shot gamble offered by lottery provider j, such that, after consumers choose between casinos, lotteries, and doing nothing, all casinos and all lottery providers earn zero average profits; and such that there are no profitable deviations from equilibrium. Specifically, there is no basic bet win probability $p_i' \neq p_i$ ($\widetilde{L}_j' \neq \widetilde{L}_j$) that casino i (lottery provider j) can offer and earn positive average profits.

We now show that there is a competitive equilibrium in which all lottery providers offer the same lottery $\widetilde{L}$ and all casinos offer the same win probability p and in which lottery providers attract the no-commitment sophisticates while casinos attract the naive agents and the commitment-aided sophisticates. To construct such an equilibrium, it is sufficient to find a lottery $\widetilde{L}$ that solves

$$\max \; V(\widetilde{L}) \tag{22}$$

– in words, $\widetilde{L}$ has the highest possible cumulative prospect theory value $V(\widetilde{L})$ among all one-shot lotteries that can be dynamically constructed, using some exit strategy, from a hypothetical casino with T rounds of gambling and a basic bet of $(\$h, q; -\$h, 1-q)$ for some $q \in (0, 0.5]$ – subject to the zero profit condition for lottery providers,

$$-\mu_{S,NC} E(\widetilde{L}) = \mu_{S,NC} C, \tag{23}$$

the participation constraint $V(\widetilde{L}) \geq 0$, and the incentive compatibility constraint, namely that the no-commitment sophisticates prefer $\widetilde{L}$ to a casino with a basic bet win probability of p; and a $p \in (0, 0.5]$ that solves

$$\max_{s \in S_{(0,1)}} V(\widetilde{G}_s) \tag{24}$$

– in words, it is the value of p that, in a casino with a basic bet win probability of p, allows agents to dynamically construct a gamble with the highest possible cumulative prospect theory value – subject to the zero profit condition for casinos,

$$-\mu_N E(\widetilde{G}_N) - \mu_{S,CA} E(\widetilde{G}_{S,CA}) = (\mu_N + \mu_{S,CA}) C, \tag{25}$$

where $\widetilde{G}_N$ and $\widetilde{G}_{S,CA}$ are random variables that measure the accumulated gains and losses under the naive agent's exit strategy and the commitment-aided sophisticate's exit strategy, respectively, and subject to the participation constraints and incentive compatibility constraints for both naive agents and commitment-aided sophisticates. If we can find such $\widetilde{L}$ and p, then there is an equilibrium in which all lottery providers offer $\widetilde{L}$ and all casinos offer a basic bet win probability of p. In particular, by construction of $\widetilde{L}$ and p, there are no profitable deviations for either casinos or lottery providers.[30]

We now construct an equilibrium explicitly. We find that the intuition underlying our equilibrium is robust, in that we are able to construct an equilibrium of the form described above for a wide range of model parameters.

Suppose that, as in Section 3, $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$, $T = 5$, and $h = \$10$; and also that $(\mu_N, \mu_{S,NC}, \mu_{S,CA}) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $C = 2$. Then there is an equilibrium in which each lottery provider offers consumers the one-shot positively skewed gamble

$$(\$50, 0.018; \$30, 0.068; \$10, 0.056; \$0, 0.309; -\$10, 0.550). \tag{26}$$

This lottery solves problem (22) subject to the associated conditions. Its expected value is -2, its cumulative prospect theory value is 1.78, and it can be dynamically constructed from a hypothetical casino offering a basic bet of $(\$10, 0.45; -\$10, 0.55)$, so that $q_j = 0.45$ for all lottery providers. Meanwhile, each casino offers the basic bet $(\$10, 0.465; -\$10, 0.535)$; in particular, $p = 0.465$ solves problem (24) subject to the associated conditions. The distribution of accumulated gains and losses with the highest cumulative prospect theory value that can be constructed out of this casino is

$$(\$50, 0.022; \$30, 0.075; \$10, 0.058; \$0, 0.311; -\$10, 0.535). \tag{27}$$

This gamble has an expected value of -1.43 and a cumulative prospect theory value of 2.15.[31]

Note that, in this equilibrium, the no-commitment sophisticates do indeed prefer the one-shot gamble (26) offered by the lottery providers to any casino. The lottery has positive cumulative prospect theory value. If these agents played in a casino, their time inconsistency would generate a negatively skewed, and hence unattractive, distribution of accumulated gains and losses. The expected value of the lottery in (26) is exactly equal to the cost, C, thereby allowing lottery providers to break even.

The commitment-aided sophisticates, however, prefer casinos because they offer better

[30] Note that $E(\widetilde{G}_N)$ is the expected value of the accumulated gains and losses under the naive agent's *actual* exit strategy, not his planned exit strategy. Because of the agent's naivete, the two strategies are, of course, different.

[31] We only report a few significant figures so as not to clutter the exposition. More precisely, in equilibrium, $q_j = 0.4499$ for all j, and $p = 0.4652$.

odds: the basic bet in a casino has a win probability of $p = 0.465$, while the lottery in (26) corresponds to a basic bet win probability of $q = 0.45$. Put differently, in a casino, the commitment-aided sophisticates can construct the accumulated gains and losses in (27) whose prospect theory value of 2.15 is higher than the 1.78 prospect theory value of the lottery in (26).

The naive agents also prefer casinos because they *think* that, in a casino, they can dynamically construct the gamble in (27), a gamble with higher prospect theory value than the lottery in (26). However, because of their time inconsistency, their actual exit strategy is quite different from their planned exit strategy. In particular, they gamble for longer in the casino, on average, than they were expecting to. As a result, the expected value of their accumulated gains and losses under their *actual* exit strategy, namely -2.57, is much lower than the expected value of their accumulated gains and losses under their planned exit strategy, namely -1.43. Since

$$-\frac{1}{3}(-2.57) - \frac{1}{3}(-1.43) = \frac{2}{3}(2),$$

the zero profit condition (25) for casinos is satisfied. Intuitively, casinos lose money on the commitment-aided sophisticates but make these losses up on the naive agents who gamble longer at casinos, on average, than they were planning to.

In summary, then, we have shown that casinos can survive in an economy with competitive provision of both lotteries and casinos. In equilibrium, lottery providers attract the no-commitment sophisticates. Casinos offer slightly better odds, and attract the naive agents and the commitment-aided sophisticates. They lose money on the commitment-aided sophisticates but make these losses up on the naive agents.

7.3 Analytical results for $T = 2$

Proposition 2 below presents analytical solutions to the decision problems faced by the three types of agents when $T = 2$. The proposition shows that the agents' behavior depends on whether the parameter

$$K = \frac{2^{\alpha+1-\delta+\frac{1}{\delta}}}{(1+3^{\delta})^{\frac{1}{\delta}}} = \frac{2^{a} w(\frac{1}{4})}{w(\frac{1}{2})} \tag{28}$$

is greater than λ, is between $1/\lambda$ and λ in value, or is less than $1/\lambda$.

Proposition 2: Suppose that a casino offers $T = 2$ rounds of gambling and a basic bet of $(\$h, \frac{1}{2}; -\$h, \frac{1}{2})$. Consider an agent who maximizes the cumulative prospect theory value of his accumulated winnings at the moment he leaves the casino and whose preference parameters satisfy $\alpha, \delta \in (0,1)$ and $\lambda > 1$. Then the agent's behavior depends on whether he is naive, a no-commitment sophisticate, or a commitment-aided sophisticate, in the following way:

Naive agent:

If $K > \lambda$, he enters the casino with a plan to stop after a first-round loss and to continue gambling after a first-round win. However, he is time inconsistent: after a first-round loss, he continues gambling and after a first-round win, he stops.

If $1/\lambda \leq K \leq \lambda$, he does not enter the casino.

If $K < 1/\lambda$, he enters the casino with a plan to stop after a first-round win and to continue gambling after a first-round loss. Moreover, he is time consistent and follows through on this plan.

No-commitment sophisticate:

If $K \geq 1/\lambda$, he does not enter the casino.

If $K < 1/\lambda$, he enters the casino with a plan to stop after a first-round win and to continue gambling after a first-round loss. He indeed follows through on this plan.

Commitment-aided sophisticate:

If $K > \lambda$, he enters the casino with a plan to stop after a first-round loss and to continue gambling after a first-round win. With the help of a commitment device, he follows through on this plan.

If $1/\lambda \leq K \leq \lambda$, he does not enter the casino.

If $K < 1/\lambda$, he enters the casino with a plan to stop after a first-round win and to continue gambling after a first-round loss. Even in the absence of a commitment device, he is able to follow through on this plan.

Proof:

We first analyze the behavior of the naive agent. At time 0, this agent can choose one of the following five plans: (a) do not enter the casino at all; (b) enter the casino and exit at time $T = 1$; (c) enter the casino and exit at time $T = 2$; (d) enter the casino and exit at time $T = 2$ or in node $(1, 1)$, whichever comes first; (e) enter the casino and exit at time $T = 2$ or in node $(1, 2)$, whichever comes first.

Note that plans (b) and (c) can never be optimal because they are both dominated by plan (a). In particular, plan (a) is strictly preferred to plan (b) if

$$0 > w(\frac{1}{2})v(h) + w(\frac{1}{2})v(-h).$$

Since $\lambda > 1$, this condition always holds. Similarly, plan (a) is strictly preferred to plan (c) if

$$0 > w(\frac{1}{4})v(2h) + w(\frac{1}{4})v(-2h).$$

Since $\lambda > 1$, this condition also always holds.

To determine the naive agent's optimal plan at time 0, then, we need only compare plans (a), (d), and (e). Plan (d) is strictly preferred to plan (a) if

$$w(\frac{1}{2})v(h) + w(\frac{1}{4})v(-2h) > 0,$$

which is equivalent to the condition $K < 1/\lambda$. Plan (e) is strictly preferred to plan (a) if

$$w(\frac{1}{4})v(2h) + w(\frac{1}{2})v(-h) > 0,$$

which is equivalent to the condition $K > \lambda$. Finally, plan (e) is preferred to plan (d) if

$$w(\frac{1}{4})v(2h) + w(\frac{1}{2})v\,(-h) > w(\frac{1}{2})v(h) + w(\frac{1}{4})v(-2h),$$

which is equivalent to the condition $K > 1$. Since $\lambda > 1$, this means that, if $K > \lambda$, the naive agent selects plan (e); that, if $1/\lambda \leq K \leq \lambda$, he selects plan (a); and that, if $K < 1/\lambda$, he selects plan (d).

Now suppose that the naive agent enters the casino. What does he actually do thereafter? Note that, if he arrives in node $(1, 1)$, he always stops because

$$v(h) > w(\frac{1}{2})v(2h) \tag{29}$$

for all $\alpha, \delta \in (0, 1)$. If, on the other hand, he arrives in node $(1, 2)$, he always continues to gamble because

$$v(-h) < w(\frac{1}{2})v(-2h) \tag{30}$$

for all $\alpha, \delta \in (0, 1)$.

In summary, then, when $K > \lambda$, the naive agent enters the casino with plan (e) in mind. However, he is time inconsistent: after he enters, he actually follows plan (d). When $K < 1/\lambda$, he enters the casino with plan (d) in mind. Moreover, he is time consistent: after he enters, he indeed follows plan (d). When $1/\lambda \leq K \leq \lambda$, he does not enter the casino at all.

Given these results, we can immediately summarize the behavior of the commitment-aided sophisticate. If $K > \lambda$, this agent enters the casino with plan (e) in mind and, with the help of a commitment device, is able to follow through on this plan. If $K < 1/\lambda$, he enters the casino with plan (d) in mind and, even in the absence of a commitment device, follows through on this plan. If $1/\lambda \leq K \leq \lambda$, he does not enter the casino at all.

To complete the proof, we consider the case of the no-commitment sophisticate. Given that conditions (29) and (30) hold for all $\alpha, \delta \in (0, 1)$, this agent knows that he will stop gambling if he arrives in node $(1, 1)$ and that he will continue to gamble if he arrives in node $(1, 2)$. He therefore knows that, if he enters the casino, the distribution of his accumulated winnings at the moment he exits the casino will be

$$(\$h, \frac{1}{2}; \$0, \frac{1}{4}; -\$2h, \frac{1}{4}).$$

He therefore enters the casino if

$$w(\frac{1}{2})v(h) + w(\frac{1}{4})v(-2h) > 0,$$

which is equivalent to the condition $K < 1/\lambda$. For $K < 1/\lambda$, then, the no-commitment sophisticate enters the casino with plan (d) in mind and then follows through on this plan. For $K \geq 1/\lambda$, he does not enter the casino.

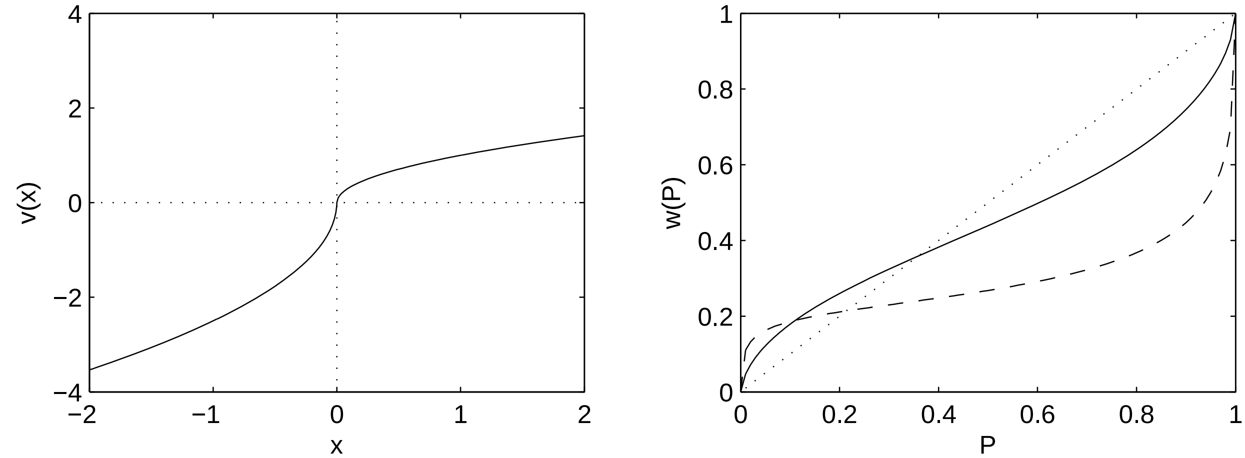

Figure 1. The left panel plots the value function proposed by Tversky and Kahneman (1992) as part of their cumulative prospect theory, namely $v(x) = x^{\alpha}$ for $x \geq 0$ and $v(x) = -\lambda(-x)^{\alpha}$ for $x < 0$, for $\alpha = 0.5$ and $\lambda = 2.5$. The right panel plots the probability weighting function they propose, namely $w(P) = P^{\delta}/(P^{\delta} + (1 - P)^{\delta})^{1/\delta}$, for three different values of δ. The dashed line corresponds to $\delta = 0.4$, the solid line to $\delta = 0.65$, and the dotted line to $\delta = 1$.

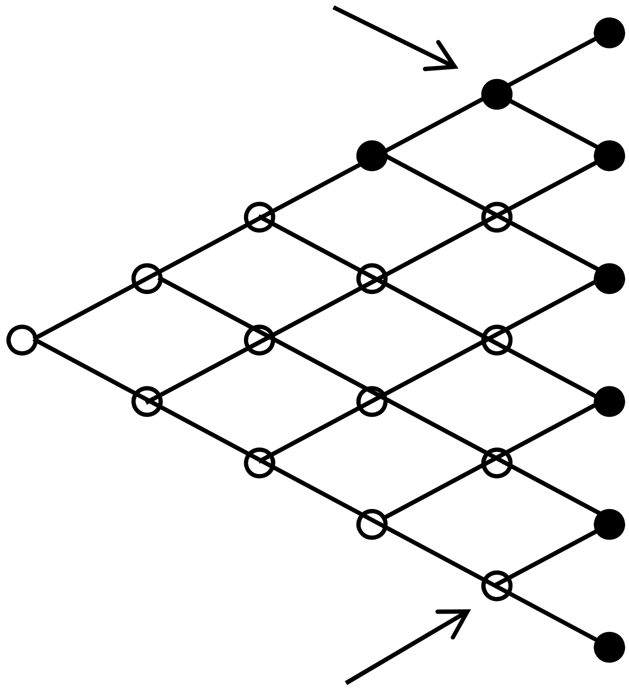

Figure 2. The figure shows how a casino can be represented as a binomial tree. Each column of nodes corresponds to a particular moment in time. The various nodes within any given column correspond to the different possible accumulated winnings or losses at that time. If a node has a black color, then the agent does not gamble at that node. At the remaining nodes, he does gamble. The arrows pick out two specific nodes that we refer to in the main text.

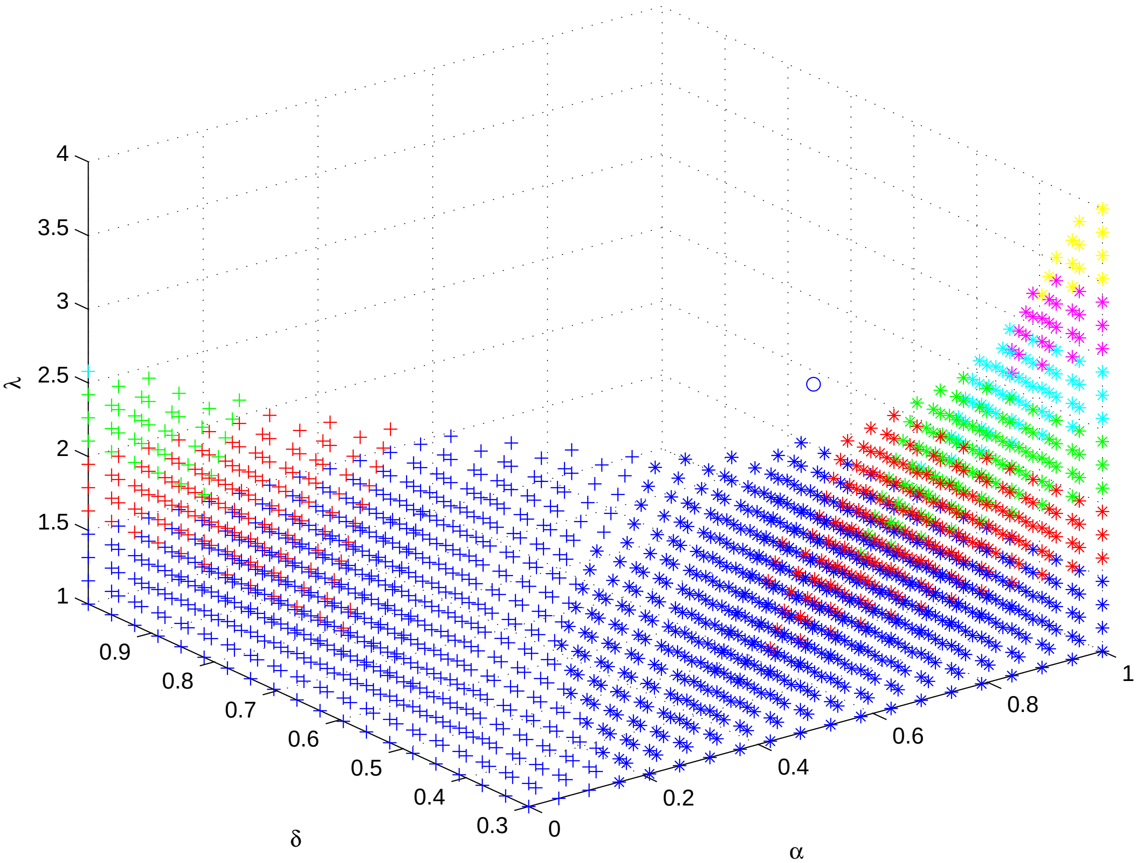

Figure 3. The "+" and "*" signs mark the preference parameter triples $(\alpha, \delta, \lambda)$ for which an agent with prospect theory preferences would be willing to enter a casino offering 50:50 bets to win or lose a fixed amount. The agent is naive: he is not aware of the time inconsistency generated by probability weighting. The "+" signs mark parameter triples for which the agent's planned strategy is to leave early if he is winning but to stay longer if he is losing. The "*" signs mark parameter triples for which the agent's planned strategy is to leave early if he is losing but to stay longer if he is winning. The blue, red, green, cyan, magenta, and yellow colors correspond to parameter triples for which λ lies in the intervals [1,1.5), [1.5,2), [2,2.5), [2.5,3), [3,3.5), and [3.5,4], respectively. The circle marks Tversky and Kahneman's (1992) median estimates of the parameters, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$. A lower value of α means greater concavity (convexity) of the prospect theory value function over gains (losses); a lower δ means more overweighting of tail probabilities; and a higher λ means greater loss aversion.

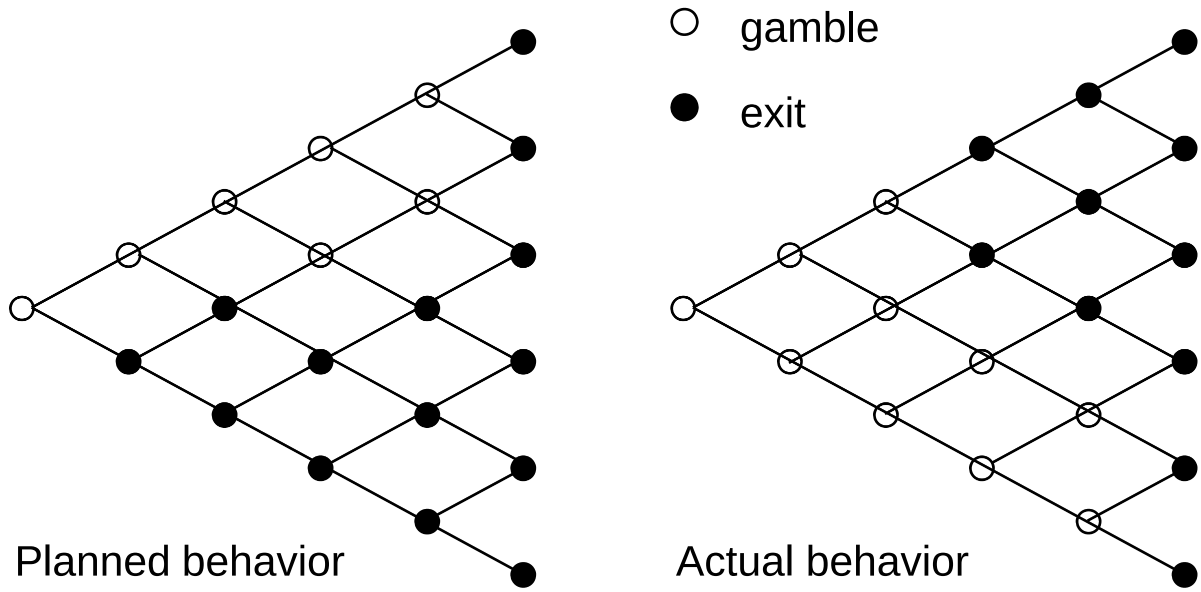

Figure 4. The left panel shows the strategy that a prospect theory agent with the preference parameter values $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$ plans to use when he enters a casino offering 50:50 bets to win or lose a fixed amount. The agent is naive: he is not aware of the time inconsistency generated by probability weighting. If a node has a black color, then the agent plans not to gamble at that node. At the remaining nodes, he plans to gamble. The right panel shows the agent's actual behavior. If a node has a black color, then the agent does not gamble at that node. At the remaining nodes, he does gamble.

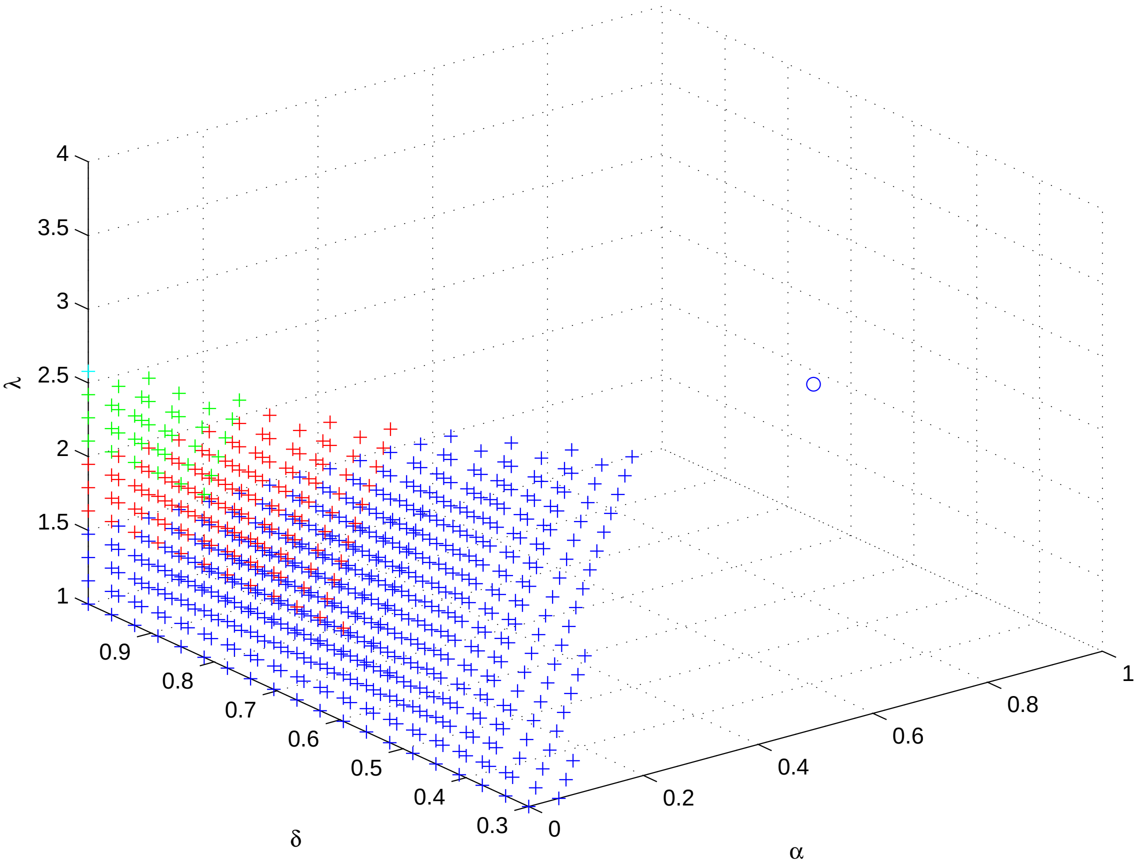

Figure 5. The "+" signs mark the preference parameter triples $(\alpha, \delta, \lambda)$ for which an agent with prospect theory preferences would be willing to enter a casino offering 50:50 bets to win or lose a fixed amount. The agent is sophisticated: he is aware of the time inconsistency generated by probability weighting. The blue, red, green, and cyan colors correspond to parameter triples for which λ lies in the intervals [1,1.5), [1.5,2), [2,2.5), and [2.5,4], respectively. The circle marks Tversky and Kahneman's (1992) median estimates of the parameters, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$. A lower value of α means greater concavity (convexity) of the prospect theory value function over gains (losses); a lower δ means more overweighting of tail probabilities; and a higher λ means greater loss aversion.