A Model of Casino Gambling

Nicholas Barberis

Yale University

March 2009*

Abstract

Casino gambling is a hugely popular activity around the world, but there are still very few models of why people go to casinos or of how they behave when they get there. In this paper, we show that prospect theory can offer a surprisingly rich theory of gambling, one that captures many features of actual gambling behavior. First, we demonstrate that, for a wide range of parameter values, a prospect theory agent would be willing to gamble in a casino, even if the casino only offers bets with zero or negative expected value. Second, we show that prospect theory 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 enters, 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 is able to commit, in advance, to his initial plan of action.

JEL classification: D03

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

*I am grateful to Daniel Benjamin, Ted O'Donoghue, Matthew Rabin, Richard Thaler, and seminar participants at 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, Yale University, and the NBER for helpful discussions. Comments are welcome at nick.barberis@yale.edu.

1 Introduction

Casino gambling is a hugely popular activity. The American Gaming Association reports that, in 2007, 54 million people made 376 million trips to casinos in the United States alone. U.S. casino revenues that year totalled almost $60 billion.

In order 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. The standard economic model of risk attitudes couples the expected utility framework with a concave utility function. 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 hard to reconcile with the standard 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. 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 (see Conlisk (1993) for a model of this last idea). A third approach suggests that gamblers simply overestimate their ability to predict the outcome of a bet; in short, they think that the odds are more favorable than they actually are.

In this paper, we present a new model of casino gambling based on Tversky and Kahneman's (1992) cumulative prospect theory. Cumulative prospect theory, one of the most prominent theories 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 posits that people 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; it is simply a modeling device for capturing the common preference for a lottery-like, or positively skewed, wealth distribution.

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 suggests that gambling is not necessarily an isolated phenomenon requiring its own unique explanation, but 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. Casinos, however, offer gambles that, aside from their low expected values, are also much less skewed than a lottery ticket. Since prospect theory agents are 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. Indeed, it has long been thought that casino gambling is the one major risk-taking phenomenon that prospect theory is *not* well-suited to explain.

In this paper, we show that, in fact, prospect theory can offer a rich theory of casino gambling, one that captures many 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 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 enters, he wants to switch to a different strategy. How a gambler behaves therefore depends on whether he is aware of this 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 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 utility 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 to gamble as long as possible if he is winning, but to stop gambling 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 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 clever 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. Through the probability weighting function, it also predicts a time inconsistency. At the moment he enters a casino, the agent's preferred plan is to keep gambling if he is winning but to stop gambling if he starts accumulating losses. We show, however, that once he starts gambling, he wants to do the opposite: to keep gambling if he is losing and to stop gambling if he accumulates a significant gain.

As a result of this time inconsistency, our model predicts significant 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 that he will exhibit a time inconsistency. This gambler *plans* to keep gambling as long as possible if he is winning and to exit only if he starts accumulating losses. After entering the casino, however, he deviates from this plan and instead gambles as long as possible when he is losing and stops only after making some gains.

The second type of agent is "sophisticated" but unable to commit: he recognizes that, if he enters the casino, he will deviate from his initial plan; but he 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 when he is losing and will stop gambling after making some gains, a strategy that will give his overall casino experience a *negatively* skewed distribution. Since he overweights the tails of probability distributions, he finds this unattractive and therefore refuses to enter the casino in the first place.

The third type of agent is sophisticated and able to commit: he also recognizes that, if he enters the casino, he will want to deviate from his initial plan; but he is able to find a way of committing to his initial plan. Just like the naive agent then, this agent plans, on entering the casino, to keep gambling as long as possible when winning and to exit only 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. According to our model, we should observe some actual gamblers behaving in this way. Anecdotally, at least, some gamblers do use techniques of this kind.

In summary, under the view proposed in this paper, casinos are popular because they cater to two aspects of our psychological make-up. First, they cater to the tendency to overweight the tails of distributions, which makes even the small chance of a large win at the casino seem very alluring. And second, they cater to what we could call "naivete," namely the failure to recognize that, after entering a casino, we may deviate from our initial plan of action.

According to the framework we present in this paper, people go to casinos because they think that, through a particular choice of exit strategy, they can give their overall casino experience a positively skewed distribution. How, then, do casinos manage to compete with another, perhaps more convenient source of positive skewness, namely one-shot lotteries? In Section 4 and in the Appendix, we use a simple equilibrium model to show that, in fact, casinos and lotteries *can* coexist in a competitive economy. In the equilibrium we describe, lottery providers attract the sophisticated agents who are unable to commit, casinos attract the naive agents and the sophisticated agents who are able to commit, and all casinos and lottery providers break even. In particular, while the casinos lose money on the sophisticated agents who are able to commit, they make these losses up by exploiting the time inconsistency of the naive agents.

Our model is a complement to existing theories of gambling, not a replacement. In particular, 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: for example, the tendency to gamble longer than planned in the region of losses, the strategy of leaving one's ATM card at home, and casinos' practice of issuing free vouchers to people who are winning.[1]

In recent years, there has been a surge of interest in the time inconsistency that stems from hyperbolic discounting.[2] While it has long been understood that probability weighting can also lead to a time inconsistency, there is very little research linking this idea to real-

[1]It is straightforward 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.

[2]See, for example, Laibson (1997), O'Donoghue and Rabin (1999), Della Vigna and Malmendier (2006), and the references therein.

world applications. In this paper, we not only analyze this second type of inconsistency in detail, but also make a case that it may be important in practice. While casino gambling is its most obvious application, it may also play a significant role in other contexts. For example, in the conclusion, we briefly mention an application to stock market trading.

In Section 2, we review the elements of cumulative prospect theory. In Section 3, we present a model of casino gambling. Section 4 discusses the model further and Section 5 concludes.

2 Cumulative Prospect Theory

In this section, we describe cumulative prospect theory. Readers who are already familiar with this theory may prefer to jump 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. Under cumulative prospect theory, the agent assigns the gamble the value

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

where[3]

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

and where $v(\cdot)$ and $w(\cdot)$ are known as the value function and the probability weighting

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

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.[4] 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 greater 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] 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 a positive 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 a negative 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.[5]

The main effect of the probability weighting in (4) is to make the agent overweight 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 in 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 both lotteries and insurance. For example, subjects typically prefer $(\$5000, 0.001)$ over a certain \$5, but also prefer a certain loss of \$5 over $(-\$5000, 0.001)$. 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 $\delta \in (0.28, 1)$.

The transformed probabilities in (3)-(4) should not be thought of as beliefs, but as decision weights which help us capture the experimental evidence on risk attitudes. In Tverksy and Kahneman's framework, an agent evaluating the lottery-like $(\$5000, 0.001)$ gamble un-

[5]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. This corrects a weakness of the original prospect theory, namely that it does not satisfy this property.

derstands that he will only receive the \$5000 with probability 0.001. The overweighting of 0.001 introduced by cumulative prospect theory is simply a modeling device which captures the agent's preference for the lottery over a certain \$5.

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 each type of gambling at \$59 billion, \$24 billion, \$4 billion, and \$200 million, respectively.[6]

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 the payoff of a casino game is typically much less positively skewed than that of a lottery ticket. And it differs from racetrack-betting and sports-betting in that casino games usually require less skill: while some casino games have an element of skill, many are purely games of chance.

In this paper, we focus our attention on casino gambling, largely because, from the perspective of prospect theory, it is particularly hard to explain. The buying of lottery tickets is already directly captured by prospect theory through the overweighting of tail probabilities. Casino games are much less positively skewed than a lottery ticket, however. It is therefore not at all clear that we can use the overweighting of tails to explain the popularity of casinos.

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.

[6]The \$200 million figure refers to sports-betting through *legal* bookmakers. It it 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.

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 at a casino.

By assuming an exogeneous date, date T, at which the agent must leave the casino if he has not already done so, we make our model somewhat easier to solve. This is not, however, the reason we impose the assumption. Rather, we impose it because we think that it makes the model more realistic: whether because of fatigue or because of work and family commitments, most people simply cannot stay in a casino indefinitely.

Of the major casino games, our model most closely resembles blackjack: under optimal play, the odds of winning a round of blackjack are close to 0.5, which matches the 50:50 bet offered by our casino. 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 the model 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 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 within a column therefore represent the different possible accumulated winnings or losses at that time.

We refer to the nodes 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, can range from 1 to $t+1$, indicates how far down the node is within the column of $t + 1$ nodes for that time: 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 colored 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 he has already left the casino in an earlier round and therefore never even reaches the 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. Clearly, a node that can only be reached by passing through a black node must itself be black. In Figure 2, the fact that node $(3, 1)$ has a black color immediately implies that node $(4, 1)$ must also have a black color.

As noted above, the basic gamble offered by the casino in our model is a 50:50 bet to 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. In fact, our analysis can easily be extended to the case in which the probability of winning $\$h$ is different from 0.5. Indeed, we find that the results we obtain below continue to hold even if, as in actual casinos, the basic gamble has a somewhat *negative* expected value: even if it entails a 0.46 chance of winning $\$h$, say, and a 0.54 chance of losing $\$h$. We discuss this issue again in Section 4.1.

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 the agent in our model *maximizes the cumulative prospect theory utility 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 making this assumption, we recognize that we are almost certainly leaving out other factors that also affect the agent's decision-making. Nonetheless, we hope to show in this and subsequent sections that our assumption is not only parsimonious but also leads to a rich theory of gambling.

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 – introduces 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 $(4, 1)$ – and ignore the specific black or white node colorations. We will see later that, *from the perspective of time 0*, the agent's preferred plan, conditional on entering the casino at all, 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 winning the \$50 prize in node $(5, 1)$. From the perspective of time 0, this prize has low probability, namely $\frac{1}{32}$, but under cumulative prospect theory, this low tail probability is 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 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 accumulated gain of \$40. If he continues gambling, he has a 0.5 chance of an accumulated gain of \$50 and a 0.5 chance of an accumulated gain of \$30. He therefore leaves the casino in node $(4, 1)$ if

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

in words, if the cumulative prospect theory utility of leaving exceeds the cumulative prospect theory utility of staying. Condition (7) simplifies to

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

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 $(4, 1)$, node $(5, 1)$ can be reached with probability 0.5. The probability 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 $(4, 1)$.

The time inconsistency in the upper part of the tree, then, is that, while the agent plans to keep gambling after accumulating some gains, he instead, if he actually makes some gains, stops gambling. There is an analogous and potentially more important time inconsistency in the *bottom* part of the tree: we will see later that, while the agent's initial plan, conditional on entering the casino at all, is to stop gambling after accumulating a loss, he instead, if he actually accumulates a loss, continues to gamble. For example, from the perspective of time 0, the agent would like to stop gambling if he were to arrive at node $(4, 5)$, the node indicated by an arrow in the bottom part of the tree in Figure 2. 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 time inconsistency. An agent who *is* aware of the time 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 related literature on hyperbolic discounting.

The first type of agent is "naive". An agent of this type does not realize that, at time $t > 0$, he will deviate from his initial plan. We analyze his behavior in Section 3.1.

The second type of agent is "sophisticated" but unable to commit. An agent of this type recognizes that, at time $t > 0$, he will deviate from his initial plan. He would therefore like to commit to his initial plan – but is unable to find a way to do so. We analyze his behavior in Section 3.2.

The third and final type of agent is sophisticated and able to commit. An agent of this type also recognizes that, at time $t > 0$, he will want to deviate from his initial plan. However, he is able to find a way of committing to this initial plan. We analyze his behavior in Section 3.3.[7]

3.1 Case I: The naive agent

We analyze the naive agent's behavior in two steps. First, we study his behavior at time 0 as he decides whether to enter the casino. If we find that, for some parameter values, he is willing to enter the casino, we then look, for those parameter values, at his behavior *after* entering the casino, 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 every node in the binomial tree between $t = 1$ 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; and "continue," which indicates that he plans to keep gambling if he arrives at that node. We denote the set of all possible plans as $S_{(0,1)}$, with the subscript $(0, 1)$ indicating that this is the set of plans that is available at node $(0, 1)$, the left-most node in the tree. Even for low values of T, the number of possible plans is very large.[8]

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 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:

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

[7]In his classic analysis of non-expected utility preferences, Machina (1989) identifies three kinds of non-expected utility agents: β-types, γ-types, and δ-types. These correspond to our naive agents, commitment-aided sophisticates, and no-commitment sophisticates, respectively.

[8]Since, for each of the $T(T + 1)/2 - 1$ nodes between time 1 and time $T - 1$, the agent can either exit or continue, an upper bound on the number of elements of $S_{(0,1)}$ is 2 to the power of $T(T + 1)/2 - 1$. For $T = 5$, this equals 16,384; for $T = 6$, it equals 1,048,576. The number of *distinct* plans is lower than 2 to the power of $T(T + 1)/2 - 1$, however. For example, for any $T \geq 2$, all plans that assign the action "exit" to nodes $(1, 1)$ and $(1, 2)$ are effectively the same.

where $V(\cdot)$ computes the cumulative prospect theory value of the gamble that is its argument. 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.

The non-concavity and nonlinear probability weighting embedded in $V(\cdot)$ make it very difficult to solve problem (9) analytically; indeed, the problem has no known analytical solution. However, we can solve it numerically and find that this approach allows us to draw out the economic intuition in full. Throughout the paper, we are careful to check the robustness of our conclusions by solving (9) for a wide range of preference parameter values.

The time inconsistency introduced by probability weighting means that we cannot use dynamic programming to solve the above problem. 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^* with the highest cumulative prospect theory value $V^* = V(\widetilde{G}_{s^*})$. The naive agent enters the casino – in other words, he plays a gamble at time 0 – if and only if $V^* \geq 0$.[9]

We now present some results from our numerical analysis. We set $T = 5$ and $h = \$10$. The shaded areas in Figure 3 show the range of values of the preference parameters α, δ, and λ for which the naive agent is willing to enter the casino, in other words, the range for which $V^* \geq 0$. To understand the figure, recall that, based on experimental data, Tversky and Kahneman's (1992) median estimates of the preference parameters are

$$(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25). \tag{10}$$

Each of the three panels in the figure fixes one of the three parameters at its median estimate and shows the range of the other two parameters for which the agent enters the casino. The small circles correspond to the median estimates in (10).

The key result in Figure 3 is 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

[9] Recall that the set $S_{(0,1)}$ consists only of plans that involve gambling at node $(0,1)$. The agent is therefore willing to gamble at this node if the best plan that involves gambling, plan s^*, offers higher utility than not gambling; in other words, higher utility than zero.

parameter values for which the agent *is* willing to enter the casino. Note that, for Tversky and Kahneman's median estimates in (10), the agent is not willing to enter the casino. Nonetheless, for parameter values that are not far from those in (10), he *is* willing to gamble.

To understand why, for some parameter values, the agent is willing to gamble, we study his optimal exit plan s^*. Consider the case of $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$; we find that, for these parameter values, the agent is willing to enter the casino. The left panel in Figure 4 shows the agent's optimal exit plan in this case. Recall that, if the agent arrives at a solid black node, he leaves the casino at that node; otherwise, he continues gambling. 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. Through extensive checks, we find that, for almost all the parameter values for which the naive agent is willing to enter the casino, the optimal exit strategy is similar to that in the left panel of Figure 4.

The exit plan in 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 casino. The reason is that, through his choice of exit plan, the agent is able to give his *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 winning the large jackpot $\$Th$ in the top-right node $(T, 1)$ becomes particularly 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 clever choice of exit strategy, to give his overall casino experience a positively skewed distribution, one which, with sufficient probability weighting, he finds attractive.[10]

We suspect that when actual gamblers enter a casino, they often have in mind a plan that is *broadly* similar to the one in the left panel of Figure 4 – specifically, a plan under which

[10] For a very small range of parameter values – a range in which α and λ are much lower than Tversky and Kahneman's (1992) estimates and δ much higher – the naive agent enters the casino with a different plan in mind, namely one in which he keeps gambling if he is losing and stops if he accumulates some gains. This strategy gives his perceived overall casino experience a negatively skewed distribution; but since α is so low and δ is so high, he does not find this unappealing.

they continue to gamble when they are winning but stop gambling once their accumulated losses reach *some* cutoff level. However, we also suspect that they may not have in mind the *exact* plan in Figure 4. In particular, they may be uncomfortable with a plan under which they might have to leave the casino after just one bet: it might feel silly to leave the casino so early if they have just traveled a long time to get there. Instead of solving (9), then, the agent may prefer to maximize $V(\widetilde{G}_s)$ over a *subset* of the plans in $S_{(0,1)}$, namely that subset for which the probability of leaving the casino in the first few rounds is lower than some given number. We find that the optimal plan in this case is, roughly speaking, one in which the agent leaves the casino only when his losses reach some cutoff level greater than zero. A plan of this kind may be more typical of the plans that many actual gamblers have in mind when they enter a casino.

Figure 3 shows that the agent is more likely to enter the casino for *low* values of δ, for *low* values of λ, and for *high* values of α. The intuition is straightforward. By adopting an exit plan under which he rides gains as long as possible but stops gambling once he starts accumulating losses, the agent gives his overall casino experience a positively skewed distribution. As δ falls, the agent overweights the tails of probability distributions all the more heavily. He is therefore all the more likely to find a positively skewed distribution attractive and hence all the more likely to enter the casino. As λ falls, the agent becomes less loss averse. He is therefore less scared by the potential losses he could incur at the casino and therefore more willing to enter. Finally, as α falls, the marginal utility of additional gains diminishes more rapidly. The agent is therefore less excited about the possibility of a large win and hence less likely to enter the casino.

We noted above that, due to the convexity of the value function in the region of losses and the use of transformed probabilities, it is difficult to solve problem (9) analytically. We are, however, able to derive the follow result, which states a sufficient condition for the naive agent to be willing to enter the casino. The proof is in the Appendix.

Proposition 1: For given preference parameters $(\alpha, \delta, \lambda)$ and a given number of rounds of gambling T, the naive agent is willing to enter the casino at time 0 if[11]

$$\sum_{j=1}^{T-\left[\frac{T}{2}\right]} (T+2-2j)^{\alpha} \left(w(2^{-T}\binom{T-1}{j-1}) - w(2^{-T}\binom{T-1}{j-2}) \right) \geq \lambda w(\frac{1}{2}). \tag{11}$$

To derive condition (11), we take one particular exit strategy which, from extensive numerical analysis, we know to be either optimal or close to optimal for a wide range of parameter values – roughly speaking, a strategy in which the agent keeps gambling when he is winning but stops gambling once he starts accumulating losses – and compute its cumulative prospect theory value explicitly. Condition (11) checks whether this value is positive; if it is, we know that the naive agent enters the casino. The condition is useful because it can shed light on the agent's behavior when T is high without requiring us to solve problem (9) explicitly, something which, for high values of T, is computationally very taxing.

For four different values of T, Figure 5 sets $\alpha = 0.88$ and plots the range of values of δ and λ for which condition (11) holds. We emphasize that the condition is sufficient but not necessary. If it holds, the naive agent enters the casino; but he may enter the casino even if it does not hold. Nonetheless, by comparing the top-left panels in Figures 3 and 5, both of which correspond to $T = 5$, we see that the parameter values for which condition (11) holds and the parameter values for which the naive agent actually enters the casino are very similar. In this sense, condition (11) is very accurate: it is not only sufficient but almost necessary as well.

The top-right and bottom panels in Figure 5 suggest that, as the number of rounds of gambling T goes up, the naive agent is willing to enter the casino for a wider range of preference parameter values. Intuitively, as T goes up, the agent, through a careful choice of exit strategy, can create an overall casino experience that is all the more positively skewed and therefore, for someone who overweights tails, all the more attractive.[12]

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

[12] It is easy to prove that the range of preference parameter values for which the naive agent enters the casino when $T = \tau$ is at least as large as the range for which he enters when $T = \tau + 1$. In particular, this follows from the fact that any plan that can be implemented in τ rounds of gambling can also be implemented in $\tau + 1$ rounds of gambling. Figure 5 gives us a sense of *how much* the range expands as T goes up.

Figures 3 and 5 show that, for Tversky and Kahneman's (1992) *median* estimates of α, δ, and λ, the prospect theory agent is only willing to enter the casino for high values of T; and Figure 5 suggests that even for high values of T, he is just barely willing to enter. There is a sense in which this fits with the evidence. Although 54 million people visited U.S. casinos in 2007, this still represents a minority of the U.S. population. The fact that the median U.S. resident does not gamble is consistent with the fact that, for the median values of the preference parameters, the prospect theory agent in our model often refuses to gamble. From the perspective of our model, the people who visit casinos are those with *lower* values of δ or λ than the median U.S. resident.

We noted earlier that we are dividing our analysis of the naive agent into two parts. We have just completed the first part: the analysis of the agent's time 0 decision as to whether or not to enter the casino. We now turn to the second part: the analysis of what the agent does at time $t > 0$. We know that, at time $t > 0$, the agent will depart from his initial plan. Our goal is to understand exactly how he departs from it.

Subsequent behavior

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

$$\max_{s \in S_{(t,j)}} V(\widetilde{G}_s). \tag{12}$$

Here, $S_{(t,j)}$ is the set of plans the agent could follow subsequent to time t, where, in a similar way to before, a "plan" is a mapping from every node between time $t+1$ and time $T-1$ to one of two actions: "exit," indicating that the agent plans to leave the casino if he reaches that node, and "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 accumulated winnings or losses the agent will experience if he exits the casino at the nodes specified by plan s, and $V(\widetilde{G}_s)$ is its cumulative prospect theory value. 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^* is the plan that solves problem (12), the agent gambles in node j at time t if

$$V(\widetilde{G}_{s^*}) \geq v(h(t+2-2j)), \tag{13}$$

where the right-hand side of condition (13) is the utility of leaving the casino at this node.[13]

To see how the naive agent actually behaves for $t \geq 1$, we return to the example from earlier in this section in which $T = 5$, $h = \$10$, and $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$. Recall that, for these parameter values, the naive agent is willing to enter the casino at time 0. The right panel of Figure 4 shows what the naive agent does subsequently, at time $t \geq 1$. By way of reminder, 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 keep gambling as long as possible when winning but to stop gambling once he started accumulating losses, he *actually*, roughly speaking, does the opposite: he stops gambling once he accumulates some gains and instead continues gambling as long as possible when he is losing. We find a similar pattern of behavior across all parameter values for which the naive agent is willing to enter the casino at time 0. Our model therefore captures a common intuition, namely that people often gamble more than they planned to in the region of losses.

Why does the naive agent behave in this way? Suppose that he has accumulated some gains. Whether he continues to gamble depends on two opposing forces. On the one hand, since he has accumulated gains, he is in the concave section of the value function. This induces risk aversion which, in turn, encourages him to stop gambling and to leave the casino. On the other hand, the probability weighting function encourages him to keep gambling: by continuing to gamble, he keeps alive the chance of winning a much larger amount of money; while this is a low probability event, the low probability is overweighted, making it attractive to keep gambling. As the agent approaches the end of the tree, however, the possibility of winning a large prize becomes less unlikely; it is therefore overweighted less, and continuing

[13]The formulation in (12) assumes that the agent's "reference point" for computing gains and losses is always fixed at his initial wealth at the moment he enters the casino. We know little about how reference points move over time. Our strategy is therefore to pick one simple assumption – that the reference point remains fixed – and to show that this leads to a rich model of gambling. Intuitively, a model in which the agent updates his reference point over time would have a harder time explaining casino gambling: in such a model, the agent would often be at the most risk averse point of the value function, the kink.

to gamble becomes less attractive. In other words, as the agent approaches the end of the tree, the concavity effect overwhelms the probability weighting effect and the agent stops gambling.

A similar set of opposing forces is at work in the bottom part of the binomial tree. Since, here, the agent has accumulated losses, he is in the convex part of the value function. This induces risk-seeking which encourages him to keep gambling. On the other hand, the probability weighting function encourages him to stop gambling: if he keeps gambling, he runs the risk of a large loss; while this is a low probability event, the low probability is overweighted, making gambling a less attractive option. The right panel in Figure 4 shows that, at *all* points in the lower part of the tree, the convexity effect overwhelms the probability weighting effect and the agent continues to gamble.[14]

3.2 Case II: The sophisticated agent, without commitment

In section 3.1, we considered the case of a naive agent – an agent who, at time t, does not realize that, at time $t' > t$, he will deviate from his time t plan. In Sections 3.2 and 3.3, we study sophisticated agents, in other words, agents who do recognize that they will deviate from their initial plans. A sophisticated agent has an incentive to find a commitment device that will enable him to stick 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" for short. In Section 3.3, we study the case of a sophisticated agent who *is* able to commit to his initial plan.

To determine a course of action, the no-commitment sophisticate uses dynamic programming, working leftward from the right-most column of the binomial tree. If he has not yet left the casino at time T, he must necessarily exit at that time. His value function in node j at time T – here, we mean "value function" in the dynamic programming sense rather than

[14]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. Over many repeated casino visits, the agent may learn his way out of the first kind of naivete. It may take much longer, however, for him 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.

in the prospect theory sense – is therefore

$$J_{T,j} = v(h(T + 2 - 2j)). \tag{14}$$

The agent then continues the backward iteration from $t = T - 1$ to $t = 0$ using

$$J_{t,j} = \max\{v(h(t + 2 - 2j)), V(\widetilde{G}_{t,j})\}, \tag{15}$$

where $J_{t,j}$ is the value function in node j at time t. The term before the comma on the right-hand side is the agent's utility if he leaves the casino in node j at time t. The term after the comma is the utility of continuing to gamble: specifically, it is the cumulative prospect theory value of the random variable $\widetilde{G}_{t,j}$ which measures the accumulated winnings or losses the agent will exit the casino with if he continues gambling at time t. The gamble $\widetilde{G}_{t,j}$ is determined by the exit strategy computed in earlier steps of the backward iteration. Continuing this iteration back to $t = 0$, the agent can see whether or not it is a good idea to enter the casino in the first place.

We now return to the example of Section 3.1 in which $T = 5$, $h = \$10$, and $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$. We find that, in this case, the no-commitment sophisticate chooses not to enter the casino at all. The intuition is straightforward. He realizes that, if he does enter the casino, he will leave as soon as he accumulates some gains but will keep gambling as long as possible if he is losing. This exit policy gives his overall casino experience a *negatively* skewed distribution. Recognizing this in advance, he decides not to enter the casino: since he overweights the tails of distributions, the negative skewness is unattractive.

The result that the no-commitment sophisticate refuses to enter the casino holds for a wide range of preference parameter values. Indeed, after extensive checks, we have been unable to find *any* $(\alpha, \delta, \lambda) \in (0.5, 1) \times (0.28, 0.8) \times (1.3, \infty)$ for which the no-commitment sophisticate is willing to enter the casino at time 0.[15]

[15] For a very small range of parameter values – a range in which α and λ are much lower than Tversky and Kahneman's (1992) estimates and δ much higher – the no-commitment sophisticate *is* willing to enter the casino. While he recognizes that his overall casino experience has a negatively skewed distribution, the fact that α is so low and δ so high means that he does not find this unappealing.

3.3 Case III: The sophisticated agent, with commitment

A sophisticated agent – an agent who recognizes that, at time $t > 0$, he will want to deviate from his initial plan – has an incentive to find a commitment device that will enable him to stick to his initial plan. In this section, we study the behavior of a sophisticated agent who is able to commit. 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 exactly the same problem as the naive agent, namely:

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

In particular, since the agent 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 $V^* \geq 0$.

Since the commitment-aided sophisticate and the naive agent solve exactly the same problem at time 0, they will, for given preference parameter values, choose exactly the same optimal strategy. Moreover, they will enter the casino for exactly the same range of preference parameter values. For $T = 5$ and $h = \$10$, for example, the commitment-aided sophisticate enters the casino for the parameter values indicated by the shaded areas in Figure 3. And for $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$, his optimal plan is the one in the left panel of Figure 4, a plan under which he continues to gamble when he is winning but stops gambling once he starts accumulating losses.[16]

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

[16] In the same way, the sufficient condition (11) for the naive agent to be willing to enter the casino is also a sufficient condition for the commitment-aided sophisticate to be willing to enter the casino.

plan: after he enters the casino, he continues to gamble when is losing and stops once he accumulates a significant gain.

Now that we have identified the strategy the commitment-aided sophisticate would like to commit to, the natural question is: *how* does he commit to it? 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?

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 in his pocket and to leave his ATM card at home. If he starts losing money, 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 casino gamblers will use a strategy of this kind. Anecdotally, at least, this *is* a common gambling strategy, which suggests that at least some of those who go to casinos fit the mold of our commitment-aided sophisticate.

In the upper part of the tree, it is less easy to think of a common strategy that gamblers use to solve the commitment problem, in other words, to keep gambling when they are winning even though they are tempted to go home. In a way, this is not surprising. One thing our model predicts – something that we have found to be especially true for higher values of T – is that the time inconsistency is much more severe in the *lower* part of the tree than in the upper part. By 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 at *every* node in the region of losses even though his initial plan was to gamble at 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 actions at 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 themselves 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 they originally planned when they are winning. This may explain the common practice among 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.

In this section, we have identified some important and arguably unique predictions of our framework. For example, our model predicts the common gambling strategy of bringing only a fixed amount of money to the casino; and it predicts the common casino tactic of giving free vouchers to people who are winning. These features of gambling have not been easy to understand in earlier models but emerge naturally from the one we present here. In particular, they are a direct consequence of the time inconsistency at the heart of our model.

Of all casino games, our model corresponds most closely to blackjack. Nonetheless, it may also be able to explain why another casino game, the slot machine, is as popular as it is. In our framework, an agent who enters the casino 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.

Throughout Section 3, we have focused primarily on the case of $T = 5$. We have also analyzed the case of $T = 10$ and find that the results for all three types of agents closely parallel those for $T = 5$. We do not use $T = 10$ as our benchmark case, however, because of its much greater computational demands. Our analysis of this case is available on request.

4 Discussion

In Section 3, we studied the behavior of three types of agents – naive agents, no-commitment sophisticates, and commitment-aided sophisticates. We now discuss some of the issues raised by this analysis: the relative average losses of the two groups that enter the casino, how casinos compete with lottery providers, and the new predictions of our framework.

4.1 Average losses

The analysis in Section 3 shows that the set of casino gamblers is made up of two distinct types: naive agents and commitment-aided sophisticates. Which of these two types loses more money in the casino, on average?

In the context of the model of Section 3 – a model in which the basic bet offered by the casino is a 50:50 bet to win or lose $\$h$ – the answer is straightforward. Since the basic bet has an expected value of zero, the average winnings are zero for both naive agents and commitment-aided sophisticates.

Now suppose, however, that the basic bet has a negative expected value, as in actual casinos. For example, suppose that the basic bet is now

$$(\$h, 0.49; -\$h, 0.51). \tag{17}$$

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 see which of naive agents and commitment-aided sophisticates has greater average losses, we therefore need to determine which of the two groups gambles for longer, on average. The group that gambles for longer will do worse.

For $T = 5$, $h = \$10$, and $(\alpha, \delta, \lambda) = (0.95, 0.5, 1.5)$, we compute the gambling behavior of the two types of agents when the basic bet has the form in (17). We find that the behavior of the naive agent is still that shown in the right panel in Figure 4 while the behavior of the commitment-aided sophisticate is still that shown in the left panel in Figure 4. This allows us to compute that the naive agent stays in the casino almost twice as long as the sophisticated agent, on average. His average losses are therefore almost twice as large. In this sense, the naivete of the naive agent – his failure to foresee his time inconsistency – is costly. We find that this result – that the naive agent stays in the casino longer, on average, than does the commitment-aided sophisticate – holds for almost all the preference parameter values for which the agents are willing to enter the casino.[17]

[17] In results not reported here but available on request, we find that, for higher values of T, the average length of stay in a casino for a naive agent is even longer, relative to that for a commitment-aided sophisticate.

4.2 Competition from lotteries

According to our model, people go to casinos because they think that, through a particular choice of exit strategy, they can give their overall casino experience a positively skewed distribution. How, then, 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 skewness that people are seeking.

One possible reason why, in reality, casinos are able to compete with lotteries – albeit a reason that lies outside our model – is because there *is* such a thing as "utility of gambling": a thrill, or rush, experienced in the moments before uncertainty is resolved. While a lottery ticket may be a more convenient source of skewness than a casino, it only offers a one-time rush. In a casino, however, people may experience a rush every time they play out a bet, thereby allowing the casino to compete with lottery providers.

Note that, even if we invoke a utility of gambling in this way, the prospect theory framework of Section 3 still makes it much easier to understand why people go to casinos. Given that casinos offer bets that have low expected values and that are often lacking in skewness, the evidence on loss aversion initially suggests that we would need to appeal to a *large* amount of gambling utility in order to explain why casinos exist. One of the insights of Section 3, however, is that even if the basic bet offered by a casino is unattractive to a loss averse agent, the agent can, through a particular choice of exit strategy, make his overall casino experience positively skewed and hence very attractive. We may therefore only need to appeal to a *small* amount of gambling utility in order to understand how casinos compete with lottery providers.

There is a second reason, however, why casinos can survive competition from lotteries – a reason that we can analyze using the framework of Section 3 and that does not rely on any notion of utility of gambling. We demonstrate the idea formally with the help of a simple equilibrium model, presented in detail in the Appendix. While we place this analysis in the Appendix, it is nonetheless an important element of our theory of casinos.

In this model, there is competitive provision of both one-shot lotteries and casinos, and yet 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, their time inconsistency will lead to a negatively skewed, and hence unattractive, distribution of accumulated gains and losses.

Casinos compete with lottery providers by offering slightly better 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 a distribution of accumulated gains and losses whose utility exceeds the utility offered by one-shot lotteries. The commitment-aided sophisticates are indeed able to construct such a distribution, and casinos lose money on these agents. Casinos make these losses up, however, on the naive agents, who, as we saw in Section 4.1, gamble in casinos longer, on average, than they were planning to. In this framework, then, casinos compete with lottery providers by taking advantage of the fact that naive agents gamble in casinos longer, on average, than do commitment-aided sophisticates, and, in particular, longer than they were initially planning to.

The equilibrium model in the Appendix also answers a closely related question, namely whether casinos would want to explicitly offer a one-shot version of the gamble 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 Predictions and other evidence

Researchers have not, as yet, had much success in obtaining large-scale databases on gambling behavior. While our model matches a range of anecdotal evidence on gambling – for example, the tendency to gamble longer than planned in the region of losses, the strategy of leaving one's ATM card at home, and casinos' practice of giving free vouchers to people who are winning – there is, unfortunately, little systematic evidence by which to judge our model.

Our model does, however, make a number of novel predictions – predictions that, we hope, can 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 time inconsistency, we should see a larger difference between planned and actual behavior among the less sophisticated.

Some recent *experimental* evidence gives us hope that these predictions will be confirmed in the field. Andrade and Iyer (2008) 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. Andrade and Iyer find that, consistent with our model, subjects systematically gamble *more* than planned after an early loss.

Another prediction comes from Figure 3, which shows that people are more likely to enter a casino if they have low values of δ and λ – in other words, if they overweight the tails of distributions more and if they are less loss averse. If we estimate δ and λ for casino goers – perhaps with the help of gambles like those used by Tversky and Kahneman (1992) – we should obtain lower values than for non-casino goers.[18]

5 Conclusion

In this paper, we present a dynamic model of probability weighting and use it to shed light on casino gambling: on why people go to casinos at all, and on how they behave when they get there.

Our framework can be applied in contexts other than casino gambling – in the context of stock trading, for example. If we think of the binomial tree of Section 3 as capturing not the accumulated gains and losses in a casino but rather the evolution of a stock price, we can

[18] A commonly heard term in the context of casino gambling is the "house money effect," the idea that people are more willing to take risk after winning some money than they were before. There is very little direct evidence of this effect from casinos, but Thaler and Johnson (1990) document it in an experimental setting. The naive agent in our model exhibits a house money effect, and he does so for the reason proposed by Thaler and Johnson (1990), namely that, after a gain, the agent moves away from the kink, the most risk averse point of the value function.

reinterpret the basic decision problem as that of a cumulative prospect theory investor who is thinking about how to trade a stock over time. As for the casino, there will be three types of traders – naive traders, no-commitment sophisticates, and commitment-aided sophisticates – with three different trading styles.

Such a framework leads to a interesting new idea, namely that some of the trading we observe in financial markets may be time-inconsistent – in other words, that people sometimes trade in ways they were not planning to. It also suggests that some of the trading rules used by asset management firms – for example, rules that require a position to be unwound if it falls more than 15% in value – may be commitment devices designed to implement trading plans that were optimal, ex-ante, but hard to stick to, ex-post. We plan to study these issues in future research.

6 References

Andrade E. and G. Iyer (2008), "Planned and Actual Betting in Sequential Gambles," *Journal of Marketing Research*, forthcoming.

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

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.

Della Vigna S. and U. Malmendier (2006), "Paying Not to Go to the Gym," *American Economic Review* 96, 694-719.

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

Gabaix, X. and D. Laibson (2006), "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics* 121, 505-540.

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

Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 62, 443-477.

O'Donoghue, E. and M. Rabin (1999), "Doing it Now or Later," *American Economic Review* 89, 103-124.

Thaler, R. and E. Johnson (1990), "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes in Risky Choice," *Management Science* 36, 643-660.

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

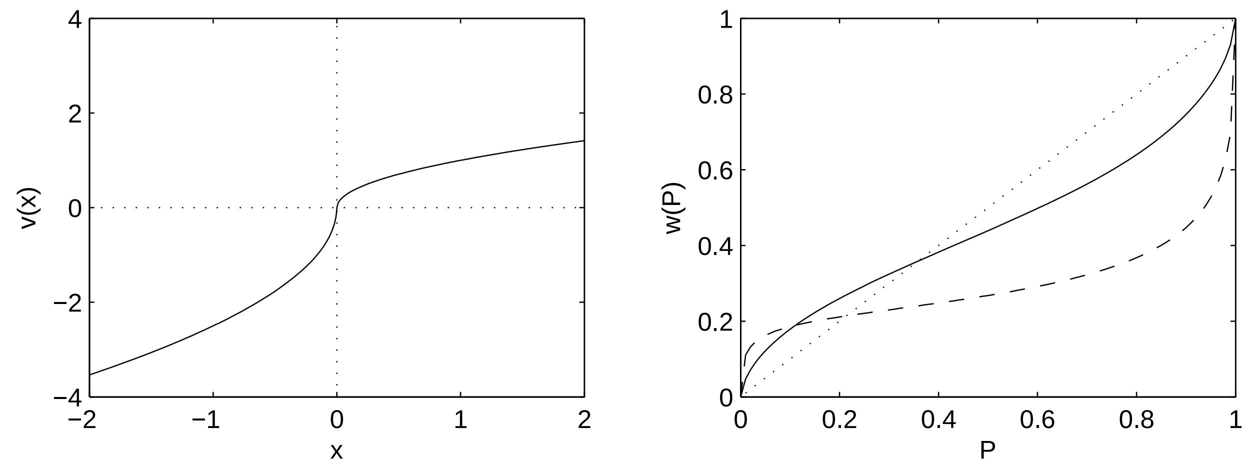

Figure 1. The left panel shows the value function proposed by Tversky and Kahneman (1992) as part of their cumulative prospect theory, namely $v(x) = x^\alpha$ for $x \geq 0$ and $v(x) = -\lambda(-x)^\alpha$ for $x < 0$, for $\alpha = 0.5$ and $\lambda = 2.5$. The right panel shows the probability weighting 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 $\delta = 1$.

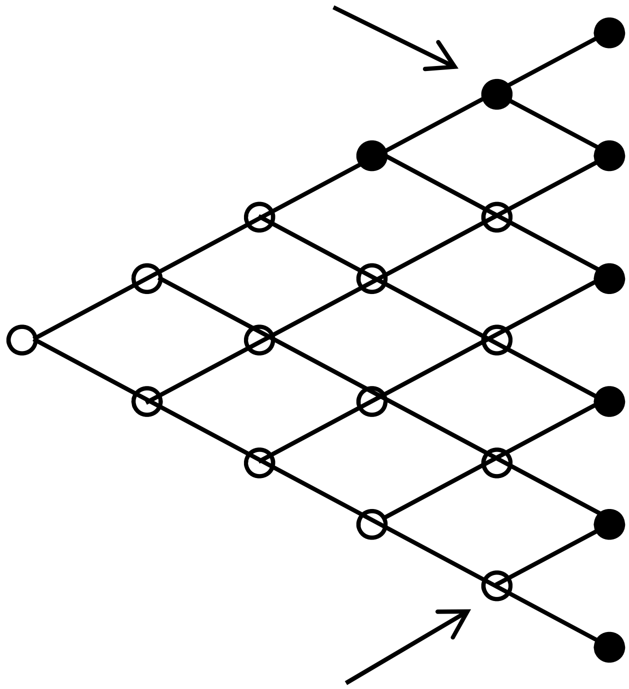

Figure 2. The figure shows how a casino can be represented as a binomial tree. Each column of nodes corresponds to a particular moment in time. Within each column, the various nodes correspond to the different possible accumulated winnings or losses at that time. A solid black node indicates that, if the agent arrives at that node, he does not gamble. At the remaining nodes, the agent does gamble.

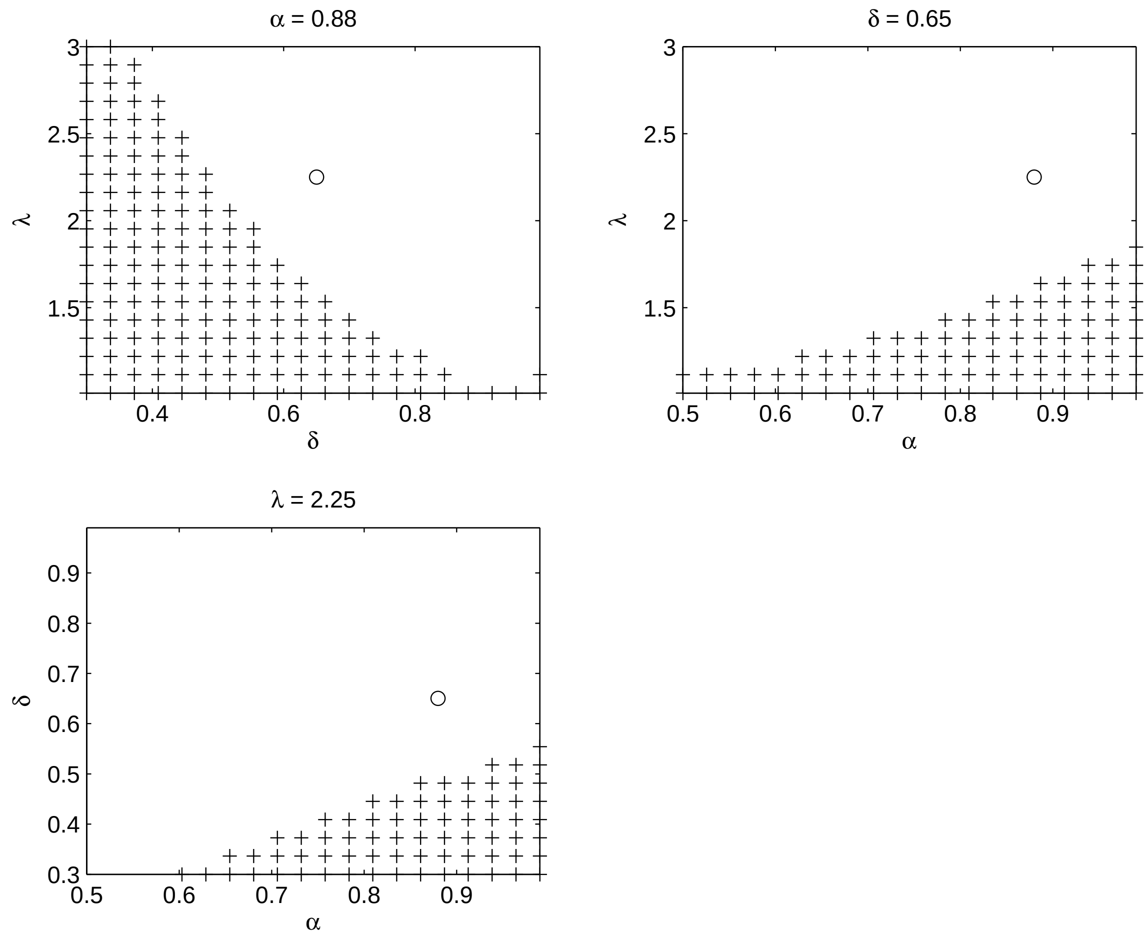

Figure 3. The "+" signs in the graphs show the range of values of the preference parameters α, δ, and λ for which an agent with prospect theory preferences would be willing to enter a casino offering 50:50 bets to win or lose a fixed amount $\$h$. The agent is naive: he does not realize that he will behave in a time-inconsistent way. Each of the three panels sets one of the three preference parameters to Tversky and Kahneman's (1992) median estimate of its value and shows the range of the other two parameters for which the agent enters the casino. The circles mark the median parameter estimates, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$.

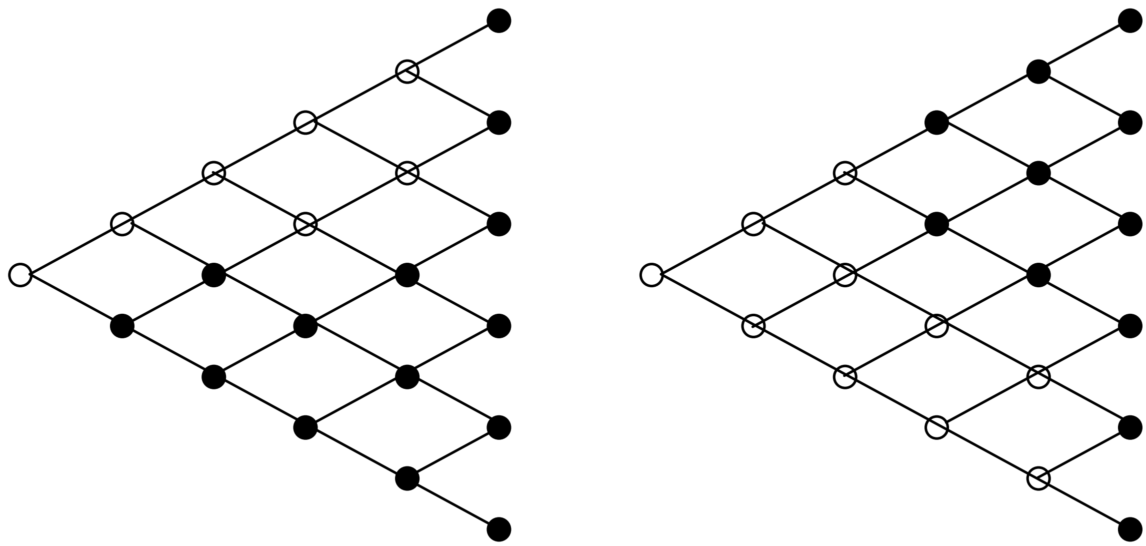

Figure 4. The left panel shows the strategy that a prospect theory agent *plans* to use when he enters a casino. The agent is naive: he does not realize that he will behave in a time-inconsistent way. If the agent arrives at a solid black node, he plans not to gamble at that node. At the remaining nodes, he plans to gamble. The right panel shows the *actual* strategy that the agent uses. If the agent arrives at a solid black node, he does not gamble. At the remaining nodes, he does gamble.

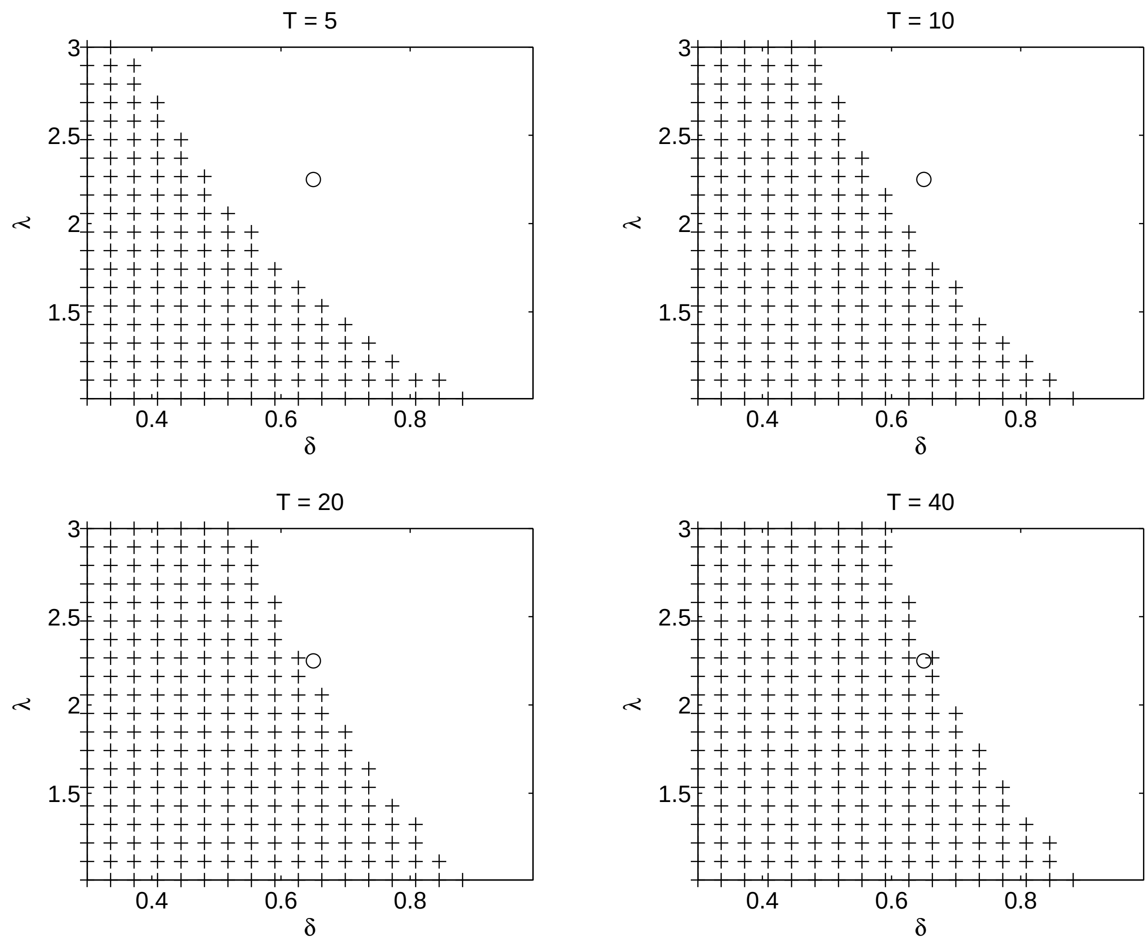

Figure 5. The "+" signs in the graphs show the range of values of the preference parameters δ and λ that satisfy a sufficient condition for an agent with prospect theory preferences to be willing to enter a casino offering 50:50 bets to win or lose a fixed amount $\$h$. The agent is naive: he does not realize that he will behave in a time-inconsistent way. The four panels correspond to four different values of T, the maximum number of rounds of gambling. In all four panels, we set the preference parameter α to 0.88. The circles mark the median parameter estimates computed by Tversky and Kahneman (1992) from experimental evidence, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$.

7 Appendix

7.1 Proof of Proposition 1

Through extensive numerical analysis, we find that when the naive agent enters the casino, he almost always 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 (11) simply checks whether the cumulative prospect theory value of this 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

$$-\lambda h^{\alpha} w(\frac{1}{2}).$$

If he exits because, at some point after the first round, his accumulated winnings equal 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 at 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 hitting 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

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

This is condition (11).

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 even if there is no explicit "utility of gambling." 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 $q_j \in (0, 0.5]$.[19]

[19] 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

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 utility. 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 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

$$\max\ V(\widetilde{L}) \tag{18}$$

– 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

but infinite utility. 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.

$q \in (0, 0.5]$ – subject to the zero profit condition for lottery providers,

$$-\mu_{S,NC} E(\widetilde{L}) = \mu_{S,NC} C, \tag{19}$$

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

$$\max_{s \in S_{(0,1)}} V(\widetilde{G}_s) \tag{20}$$

– 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,

$$-\mu_N E(\widetilde{G}_N) - \mu_{S,CA} E(\widetilde{G}_{S,CA}) = (\mu_N + \mu_{S,CA}) C, \tag{21}$$

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.[20]

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

[20]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.

lottery provider offers consumers the one-shot positively skewed gamble

$$(\$50, 0.018; \$30, 0.068; \$10, 0.056; \$0, 0.309; -\$10, 0.550). \tag{22}$$

This lottery solves problem (18) 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 (20). The distribution of accumulated gains and losses with the highest cumulative prospect theory value that can be constructed out of this casino is

$$(\$50, 0.022; \$30, 0.075; \$10, 0.058; \$0, 0.311; -\$10, 0.535). \tag{23}$$

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

Note that, in this equilibrium, the no-commitment sophisticates do indeed prefer the one-shot gamble (22) 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 (22) 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 odds: the basic bet in a casino has a win probability of $p = 0.465$, while the lottery in (22) 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 (23) whose prospect theory value of 2.15 is higher than the 1.78 prospect theory value of the lottery in (22).

The naive agents also prefer casinos because they *think* that, in a casino, they can dynamically construct the gamble in (23), a gamble with higher prospect theory value than the lottery in (22). 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 (21) 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.