

HISTORY-DEPENDENT RISK PREFERENCES: EVIDENCE FROM INDIVIDUAL CHOICES AND IMPLICATIONS FOR THE DISPOSITION EFFECT*

Angie Andrikogiannopoulou[†]
London School of Economics

Filippos Papakonstantinou[‡]
Imperial College London

August 2016

Abstract

Using trading data from a sports-wagering market, we estimate individual risk preferences within the prospect-theory paradigm. This market's experimental-like features greatly facilitate preference estimation, while our long panel enables us to study whether preferences vary across individuals and depend on earlier outcomes. Our estimates i) extend support for existing experimental findings — mild utility curvature, moderate loss aversion, probability overweighting of extreme outcomes — to a real market that shares similarities with traditional financial markets, ii) reveal that risk attitude is heterogeneous and history-dependent, and iii) indicate that prospect theory can better explain the prevalence of the disposition effect than previously thought.

*Keywords: Risk Preferences, State Dependence, History Dependence, Heterogeneity, Prospect Theory
Disposition Effect*

JEL Classification: D03, D12, D14, D81, G02, G11

*For helpful comments, we thank Nick Barberis, Kenza Benhima, Adrian Bruhin, Thomas Epper, Harald Hau, Bo Honore, Alex Kostakis, Andriy Norets, Chris Rogers, Luis Santos-Pinto, Bernard Salanie, Hyun Shin, Dimitri Vayanos, Wei Xiong, participants at the 2015 Miami Behavioral Finance Conference, and seminar participants at Athens University of Economics and Business, Cambridge University, Imperial College London Business School, London School of Economics, McGill University, Oxford University, Princeton University, Queen Mary University of London, the Stockholm School of Economics, the University of Amsterdam, the University of Geneva, the University of Lausanne, the University of Piraeus, and the Wharton School. This work has been supported by an Amazon Web Services in Education research grant, and by a Swiss Finance Institute grant.

[†]London School of Economics, Department of Finance, Houghton Street, London WC2A 2AE, UK, e-mail: a.andrikogiannopoulou@lse.ac.uk.

[‡]Imperial College Business School, Tanaka Building, South Kensington Campus, London SW7 2AZ, UK, e-mail: fpapakon@imperial.ac.uk.

1 Introduction

Even though there is a substantial body of experimental evidence supporting prospect theory, its more widespread application and acceptance in finance is hindered by the lack of non-experimental evidence on individual choice under risk. On the one hand, preference estimates from experimental settings face objections regarding their generalizability to real market settings, which involve substantial amounts of money. On the other hand, traditional financial markets do not allow for a clean identification of preference parameters, due to the complex structure of asset returns and the limited availability of trading data for individual market participants. In this paper, we use the sports wagering market as a real-world financial-market laboratory to estimate individual risk preferences within the prospect-theory paradigm.

Wagering markets possess unique features that greatly facilitate preference estimation. First, they feature a variety of assets (i.e., the wagers) that bear no systematic risk, are short-lived, and have a terminal payoff that is determined exogenously by match outcomes. Second, a wager's quoted price is a predictor of its win probability, which in a relatively efficient market (as the one we study) implies that its payoff distribution is observable *ex ante* with significant accuracy. And third, market participants typically make over time a large number of choices with serially uncorrelated payoffs. These features allow for a lottery representation of individual choices, facilitate the clean identification of utility parameters, and lead to better inference. In contrast, in traditional financial markets i) most assets do not have an exogenous settling-up point, ii) asset returns are not independent and their distribution in the future cannot be accurately estimated, and iii) individuals' trades, if observed, are infrequent. At the same time, sports wagering markets share similarities to financial markets, both in terms of their institutional organization and in terms of their participants' mentality. Thus, they provide an ideal empirical laboratory for preference estimation, that can yield valuable insights about traditional financial markets. Indeed, the literature has long identified the relevance of the sports wagering markets for traditional financial markets (e.g., Gandar et al., 1988; Gray and Gray, 1997; Avery and Chevalier, 1999; Durham, Hertz and Martin, 2005; Moskowitz, 2015). Our study differs from previous studies in that we focus on an individual-level analysis of behavior rather than an aggregate analysis based on market price data.

The contribution of this paper is threefold. First, our estimates of individual preference parameters confirm previous experimental findings in a real-world market setting: individuals are mildly risk averse (loving) over gains (losses), moderately loss averse, and overweight extreme outcomes' probabilities. Second, the richness

of our data allows for a comprehensive analysis of preference heterogeneity and history dependence, which largely exceeds the scope of existing results. Third, we explore the implications of our preference estimates for a dynamic portfolio choice problem, and we obtain new insights regarding the ability of prospect theory to explain the prevalence of the disposition effect, i.e., people's tendency to sell (hold) stocks whose value has increased (decreased) since purchase.

To get a first sense of the behavior of individuals in the wagering market, we conduct a preliminary regression analysis of the effect of past outcomes on subsequent risk-taking. We find that choices are affected by past profits, in a manner consistent with the "house-money" and "break-even" effects (Thaler and Johnson, 1990): individuals increase their betting frequencies, wager higher amounts and select lotteries with higher variance in the presence of prior gains, while after losses, they favor lotteries with significantly higher skewness, presumably because these lotteries offer them a good opportunity to break even. This behavior is at odds with the traditional expected utility theory in which the only possible channel through which past outcomes affect future choices is through a wealth effect, and naturally lends itself to a reference-dependent theory, in which past outcomes may affect the reference point hence the definition of gains and losses.

Motivated by this evidence, we model risk-taking behavior using an extension of the cumulative prospect theory (CPT) of Tversky and Kahneman (1992), which is the most prominent reference-dependent theory of choice. According to CPT, individuals i) frame choices narrowly, i.e., evaluate risks in one context separately from other risks; ii) evaluate outcomes by computing gains and losses relative to a reference point; iii) use a value function that is concave for gains and convex for losses, corresponding to risk aversion and risk seekingness, respectively; iv) have different sensitivity to gains and losses, corresponding to loss aversion; and v) systematically distort event probabilities using a probability weighting function. To incorporate history dependence in individual behavior, we extend CPT to dynamic choice situations by proposing that the reference point that separates gains from losses may not completely adapt to wealth fluctuations. In particular, inspired by the standard exponential-discounting modeling assumption in economics, we model the reference point as a function of earlier outcomes, with exponentially lower weights attached to more distant outcomes. Essentially, the discount factor is the rate of memory decay which measures how far back each individual is able to recall past outcomes. As a result, the reference point may update immediately to incorporate past profits, in which case it equals contemporaneous wealth, or it may update slowly, in which case it "sticks"

in between current and previous wealth levels, giving rise to time-varying risk preferences.

To estimate the CPT parameters, we embed our model of behavior within an econometric model that allows for unobserved preference heterogeneity across individuals. In particular, we recognize that all utility parameters are drawn from a population distribution, which we estimate from the data by looking for cross-sectional variation in individuals' choices. The variety of lotteries available in the sports wagering market allows us to identify the usual CPT parameters, and the relatively long time dimension of our panel enables us to identify the dependence of the reference point on past outcomes. Following the literature, our main identification assumption is that individuals' subjective beliefs are homogeneous and rational, as approximated by the sportbook prices which are empirically found to be close to the true probabilities of match outcomes. In a sensitivity analysis, we show that our results are robust to small but significant deviations of the subjective from the price-implied probabilities, as well as to alternative approximations of the rational beliefs. The model is estimated using a Markov Chain Monte Carlo Gibbs sampler, and its fit appears significantly better than alternative models that do not incorporate the CPT features, preference heterogeneity, and/or history dependence.

We find that, on average, the curvature of the value function is mild, loss aversion is moderate, and individuals overweight extreme positive and negative outcomes. Though some of our parameter estimates deviate from the widely used estimates from the Tversky and Kahneman (1992) experiment, they are all well within the range of estimates found in the experimental literature (see Booij, van Praag and van de Kuilen, 2010 for a meta-analysis).¹ Thus, our results validate the general features of individual risk-taking behavior that have been estimated in experiments, and show that they carry over to real-world market settings where individuals experience real gains *and* losses and face choices with a variety of payoffs and probabilities. This is significant, given that an increasing number of studies in finance appeal to these types of preferences to explain various behaviors and phenomena observed in financial markets, ranging from stock market (non-) participation and under-diversification (e.g., Polkovnichenko, 2005; Mitton and Vorkink, 2007), to the disposition effect (e.g., Kyle, Ou-Yang and Xiong, 2006; Li and Yang, 2013), to the equity premium puzzle and the pricing of idiosyncratic skewness (e.g., Barberis and Huang, 2008; Zhang, 2006).

Furthermore, we find considerable heterogeneity in risk attitude across individuals, which suggests that the

¹Tversky and Kahneman (1992) estimate, on average, a value function with a similar curvature to ours and a more pronounced kink at the reference point corresponding to stronger loss aversion, and a probability weighting function that overweights not only extreme outcomes but also outcomes near the reference point.

representative-agent assumption often used in finance applications likely misses important aspects of individual behavior. Incorporating preference heterogeneity in finance can have important implications. For example, at the individual level, it can help us understand the observed heterogeneity in households' portfolio choices, from, e.g., stock market participation to portfolio diversification (for a review, see Curcuru et al., 2010). At the aggregate level, preference heterogeneity can help explain observed patterns in trading volume, the equity premium puzzle, equilibrium asset prices, and the trading of assets with zero net supply such as derivatives (e.g., Wang, 1996; Benninga and Meishar, 2000; Chan and Kogan, 2002; Bhamra and Uppal, 2009).

Finally, our estimates reveal that risk preferences are history-dependent, as the reference point most individuals use to separate gains from losses sticks to previous levels. Notably, the estimated distribution of the memory decay parameter reveals a continuum of individuals who are affected to varying degrees by previous outcomes, with a marked concentration close to the two natural extremes of no memory and perfect memory. Incorporating history dependence, hence time-variation, in risk preferences also has important implications in finance. At the individual level, time-varying risk preferences can be useful for understanding why households' portfolio choices vary in response to wealth fluctuations and macroeconomic shocks (e.g., Massa and Simonov, 2005; Calvet, Campbell and Sodini, 2009; Guiso, Sapienza and Zingales, 2013) and for explaining patterns in dynamic trading behavior such as the disposition effect (Shefrin and Statman, 1985; Odean, 1998). At the aggregate level, they can help explain the dynamics of equilibrium asset prices and match stylized facts about asset returns such as the high mean and counter-cyclicalities of the equity premium and the excess volatility of stock returns (e.g., Constantinides, 1990; Campbell and Cochrane, 1999; Barberis, Huang and Santos, 2001).

In contrast to our estimation results, most applied behavioral models in finance assume i) a representative-agent, ii) the median CPT parameter estimates of Tversky and Kahneman (1992), and iii) that individuals have either no memory or perfect memory. Although a complete analysis of the general equilibrium asset pricing implications of our preference estimates is outside the scope of the current paper, we do consider some partial equilibrium implications on individual dynamic trading behavior, specifically on the disposition effect. Barberis and Xiong (2009) solve a dynamic portfolio-choice problem with a representative prospect-theory agent with perfect memory to show that — contrary to conventional wisdom — prospect theory does not provide a plausible explanation of the disposition effect. In particular, they show that this agent would optimally not buy equity for small values of the risk premium, while for larger values he would be more inclined to sell stocks

with prior losses than stocks with prior gains, i.e., he would exhibit the *opposite* of the disposition effect. In an application, we solve a two-period portfolio-choice problem similar to Barberis and Xiong (2009), replacing the representative-agent assumption with our heterogeneous history-dependent risk-preference estimates. Our findings provide a more encouraging message with respect to prospect theory’s relevance in explaining the prevalence of the disposition effect. We show that prospect theory *can* explain the disposition effect for a wide range of equity risk premia. In particular, for the empirically observed equity risk premium and trading frequency, more than half of our individuals would optimally buy equity and more than half of those would exhibit the disposition effect. The difference in our findings from those of Barberis and Xiong (2009) can be attributed to the fact that we estimate substantial heterogeneity across individuals in the preference parameters as well as a reference point that is sticky — but not *perfectly* sticky — for most individuals.

This paper is also related to two strands of the empirical literature: studies that structurally estimate risk preferences and studies on the dynamics of risk taking. The vast majority of structural-estimation studies use data from lab experiments or even quasi-experiments in non-market settings such as TV game shows.² These have uncovered many useful insights about individual behavior, most notably that risk preferences are widely heterogeneous across individuals and often violate expected utility theory. A small number of studies estimate preferences in the field using real-world data. This comes at a cost — loss of experimental control — but it also comes with the benefit of increased external validity, as field data capture individuals’ naturally arising behavior. The majority of these field studies use aggregate price data to elicit the preferences of a representative agent. For example, Golec and Tamarkin (1998), Jullien and Salanie (2000), and Snowberg and Wolfers (2010) use data on horse-racetrack betting prices, and Kliger and Levy (2009) and Polkovnichenko and Zhao (2013) use data on option prices. Only a handful of studies utilize data on individuals, hence are able to refrain from the representative-agent assumption and estimate heterogeneous preferences: Cicchetti and Dubin (1994), Cohen and Einav (2007) and Barseghyan et al. (2013) use individuals’ choice of deductible when buying insurance, and Paravisini, Rappoport and Ravina (2015) use data from person-to-person lending. Our paper complements the latter group of studies by analyzing individual-level choices in a sports wagering market. Our data have two distinct features relative to the data used in these studies. First, we observe an order

²For lab experiments, see, e.g., Hey and Orme (1994), Choi et al. (2007), von Gaudecker, van Soest and Wengstrom (2011). For TV game show studies see, e.g., Gertner (1993), Metrick (1995), Post et al. (2008).

of magnitude more choices for each individual, which enables us to estimate the extent to which preferences vary not only across individuals but also over time. Second, we observe a greater variety of lotteries in terms of payoffs (both positive and negative) and probability distributions, which enables us to estimate all features of prospect theory (value-function curvature, loss aversion, and probability weighting) for a wide range of prizes and probabilities.³ Andrikogiannopoulou and Papakonstantinou (2016) develop a mixture model of different utility types and use the same data to estimate preferences, focusing on heterogeneity *within* and *across* types. The main findings there are that utility curvature alone cannot explain the choices of individuals, and that probability weighting is a prevalent feature of risk attitudes that is exhibited by almost all individuals. Thus, using CPT as the starting point, in this paper we fully exploit the panel structure of the data to estimate preferences with heterogeneity *and* history dependence, and we analyze their implications for dynamic trading behavior.

In the empirical literature that studies the dynamics of individual risk taking, history dependence was first documented in the experiment of Thaler and Johnson (1990), who showed that people are more likely to take risk after gains (house-money effect), but are also more likely to take risks that give them a chance to recover prior losses (break-even effect). Similar behaviors have been observed in TV game shows by Gertner (1993) and Post et al. (2008). In the field, Odean (1998) finds that households exhibit the disposition effect, Coval and Shumway (2005) find that professional traders with morning gains (losses) make less (more) risky trades in the afternoon, while Liu et al. (2010) also document history dependence in professionals' behavior, but in the opposite direction. With the exception of Post et al. (2008), these studies conduct reduced-form regression analyses to estimate an average relationship between prior outcomes and measures of risk taking.⁴ In contrast, our structural approach allows us to estimate the underlying preference parameters that give rise to this behavior.

We note that the literature on preference estimation, including our paper, can be subject to two potential criticisms. First, each study focuses on a specific group of people (e.g., game show participants or people shopping for insurance), so the derived estimates may not be easily generalizable to other groups. Even

³Of the field studies that incorporate heterogeneity in risk preferences, Cicchetti and Dubin (1994), Cohen and Einav (2007), and Paravisini, Rappoport and Ravina (2015) focus on expected utility theory, while Barseghyan et al. (2013) incorporate one of the CPT features (probability weighting) in their model. Since, in the insurance domain, all outcomes are losses relative to initial wealth, loss aversion in its original form is not identified, and probability distortions can be estimated only for the domain of losses and only for a small range of probabilities (between 0% and 16%; see Barseghyan et al., 2013). Furthermore, the existing studies do not incorporate history dependence due to the limited number of choices per individual.

⁴Due to the nature of their data, the model of Post et al. (2008) differs from ours and the original CPT (no heterogeneity, no probability weighting, and the reference point / loss aversion are related to future expectations). Their implementation of history dependence is also specific to their setting (see Section 4.1 for a more detailed discussion).

though the group of people placing wagers on sports may be self-selected, it is actually quite large: Recent surveys in the U.S. (by ESPN, in 2008) show that about 50% of the adult population and 67% of college students place sports wagers each year, and it is estimated that close to \$1 trillion is wagered on sports, globally, per year (H2 Gambling Capital report of 2013). Indeed, the fact that the risk preferences we estimate are not only very heterogeneous but also well within the range of extant experimental estimates suggests that a potential sample selection bias is likely not severe. The second potential criticism is that each preference estimation study uses data from a specific setting, so the derived estimates may not be easily generalizable to other domains. For example, in the insurance context, individuals may be primed to make safe choices that limit their exposure, so estimates may be less relevant for, e.g., the stock market where people's primary motive is "financial gain" (Hoffman, 2007). Similarly, it could be argued that a key driver of individuals' behavior in the sports wagering market is non-pecuniary, e.g., entertainment or team fandom.

Keeping in mind that extrapolating preference parameters from *any* context to another necessitates additional assumptions, our estimates provide a valuable complement to the existing estimates in the literature, and could be particularly relevant for traditional financial markets for the following reasons. First, the sports wagering market operates in a similar way to the stock market: A large number of agents risk money on the uncertain outcomes of future events, sports bookmakers are analogous to market makers, sports handicappers play the role of financial analysts, and information about sports events is widely available in the media. Second, survey evidence shows that — like in the stock market — the primary driver of individuals' behavior in the sports wagering market is financial gain ("to make money" or to have "the chance of winning big money") and non-pecuniary motives like entertainment and team loyalty are an important but secondary motive.⁵ Conversely, recent studies show that many individuals in the stock market also view trading as an entertaining and sensation-seeking activity, are motivated by loyalty, and prefer stocks with lottery characteristics.⁶ Third, while — at first glance — it seems that individuals who participate in a market

⁵For evidence on individual behavior in sports wagering markets, see the 2010 British Gambling Prevalence Survey. Consistent with this evidence, we find that team loyalty/fandom does not play a strong role in individuals' choices in our data, as individuals place, on average, fewer than 3% of their wagers on any one team.

⁶In a survey, Hoffman (2007) finds that the second most important reason for investing is that it is "a nice free-time activity," behind investing for "financial gain" but ahead of investing to "safeguard [one's] retirement." Dorn and Sengmueller (2009) find that the vast majority of investors they survey "enjoy investing" in stocks and that entertainment motives partially drive trading; similarly, Grinblatt and Keloharju (2009) show that sensation-seeking partially drives trading. Gao Bakshi and Lin (2015) find that when there is a large jackpot lottery in Taiwan, some individuals substitute toward buying lottery tickets and away from trading stocks. Kumar (2009) shows that individuals prefer low-priced stocks with high idiosyncratic volatility and skewness.

with negative expected returns exhibit risk-seeking behavior which is at odds with the major stylized fact that stock returns are on average positive, observed behavior is actually quite consistent across the two markets. Indeed, it has been shown that people's behavior in settings that offer negative expected returns is better explained by a preference for *skewness*, not *risk* (see Golec and Tamarkin, 1998, Snowberg and Wolfers, 2010, and Garrett and Sobel, 1999 for evidence from horse-racetrack betting and state lottery games). This is also true for the individuals in our data, as the value function we estimate is generally concave, which implies an aversion to risk, while there is substantial probability weighting, which implies a preference for skewness. These preference characteristics — risk aversion coupled with a preference for skewness — are consistent with a number of stylized facts about pricing in traditional financial markets. Barberis and Huang (2008) show that CPT investors with a concave value function and probability weighting dislike variance and require a *positive* risk premium, as observed in the stock market. They also argue that these CPT preferences induce a preference for skewness and a *negative* premium for idiosyncratic skewness in equilibrium, which may explain the low returns associated with 'assets' with positively skewed returns, such as IPOs (Ritter, 1991), private business holdings (Moskowitz and Vissing-Jørgensen, 2002), distressed stocks (Campbell, Hilscher and Szilagyi, 2011), and options on individual stocks (Bollen and Whaley, 2004).⁷

The rest of the paper is structured as follows. In Section 2, we describe the data. In Section 3, we conduct a preliminary reduced-form analysis of the relationship between past profits and subsequent risk taking. In Section 4, we present our structural model of behavior and its econometric implementation, our preference estimates, robustness checks, and intuition regarding the implications of our estimates for risk taking after gains and losses. In Section 5, we analyze the implications of our estimates for the disposition effect. In Section 6, we conclude.

⁷Several studies provide direct evidence that investor behavior is consistent with these types of preferences. Mitton and Vorkink (2007) show that household portfolio under-diversification is related to a preference for skewness, Zhang (2006), Boyer, Mitton and Vorkink (2010), and Bali, Cakici and Whitelaw (2011) show that stocks with extreme positive or positively skewed returns have negative mean returns, and Ni (2008) and Boyer and Vorkink (2014) show the same for individual stock options. Lemmon and Ni (2014) show that individual stock options, unlike index options, are traded heavily by individual investors, and Conrad, Kapadia and Xing (2014) show the same for stocks with high default probability. Conrad, Kapadia and Xing (2014) calibrate the Barberis and Huang (2008) model and show that the predicted mean returns for stocks with high default probability match those in the data.

2 Data

Here, we provide information about the sports wagering market, we describe our data, and we explain how we create the lottery representation of individual bet choices.

2.1 The sports wagering market

The sports wagering market is a quote-driven dealer market that is run by the bookmaker and offers individuals the possibility to make trades relating to the outcomes of sporting events. The bookmaker quotes the current odds, i.e., the inverse of the price of a unit monetary payout, for each outcome of each event. Individuals can place a wager at these prices and the bookmaker takes the opposite side. While the bookmaker may change the prices over time, the payoff of each wager is determined by the prices prevailing at the time the wager was placed.⁸ For example, if an event has two possible outcomes with prices 0.40 and 0.65 (quoted as having odds $\frac{1}{0.40} = 2.50$ and $\frac{1}{0.65} \approx 1.53$), respectively, then an individual who backs the first (second) outcome will make a profit of 1.50 (0.53) for each 1.00 staked if he wins. Associated with these prices are the *implied probabilities* of the outcomes, defined as $\frac{0.40}{0.40+0.65} \approx 0.38$ and $\frac{0.65}{0.40+0.65} \approx 0.62$, which will be useful for our analysis later.

The bookmaker's general objective is to make money through a commission that is incorporated in the odds, much like a bid-ask spread. In the example above, it is clear that an individual needs to wager $0.40 + 0.65 = 1.05$ to receive a 1.00 payoff with certainty, in which case he loses 0.05 and the bookmaker makes 0.05 with certainty. The traditional model of sportsbook pricing suggests that the bookmaker balances the book for each event by adjusting prices so that he makes the same profit regardless of the outcome of the event. That is, if at the current prices one of the outcomes in an event is heavily bet, the bookmaker could increase its price to shift betting activity to the other outcomes so that the total payout to winners, hence his commission, would be the same regardless of the realized outcome. Alternatively, the bookmaker may deliberately set prices near the efficient ones, occasionally allowing the book to become slightly unbalanced as dictated by betting activity (for evidence consistent with this strategy, see e.g., Paul and Weinbach, 2008, 2009). This strategy saves on the costs and risks associated with perfectly balancing the book at all times, while it is not particularly risky for the bookmaker since the commission he charges provides a cushion against the unbalanced liabilities implied

⁸This describes fixed-odds betting, which is commonly used in sports wagering. In contrast, in parimutuel betting which is more commonly used in horse racetrack betting, individuals place bets by putting money in a pool and payoff odds are determined when the pool is closed and are calculated by sharing the pool among the winners.

by an unbalanced book, hence overall it could lead to greater long-run profits.⁹ For example, given an outcome that is expected to occur with probability 0.38, if the bookmaker wishes to earn a commission of 5%, he will set the outcome's price at $\frac{0.38}{1-0.05} = 0.40$ rather than the actuarially fair 0.38. That is, the outcome will be quoted as having odds $\frac{1}{0.40} = 2.50$, hence in expectation the individual will make $0.38 \cdot (+1.50) + 0.62 \cdot (-1.00) = -0.05$ and the bookmaker will make 0.05. Even though a thorough study of the bookmaker's price-setting behavior is beyond the scope of (and the data available for) this paper, an analysis of our data and discussions we had with the bookmaker who provided it to us give support to a pricing behavior that is closer to the efficient pricing model. Regardless, as we discuss in Section 4.5, our results are robust to alternative behaviors.

2.2 Data description

We use panel data of individual activity in a large European online sports wagering company. The data set contains information about the activity of 336 randomly selected customers over a 5-year period (October 2005 – November 2010).¹⁰ The sportsbook under study offers wagers on a wide range of sports, a large number of matches within each sport, and a variety of events related to each match (e.g., given a baseball match, one can place a wager on the winner, the total number of runs, etc.). An individual can place wagers either separately as single bets or together in combination bets, to produce a wide variety of payoff and risk profiles. The individual then chooses the amount of money to stake on each bet, and fills out a betting slip like the one in Figure 1, which summarizes all the relevant information.

For each bet placed by each individual in our sample, we observe the following: i) bet date, ii) bet event and outcome chosen, iii) bet amount, iv) bet type, v) prices for all outcomes of the bet event, and vi) bet result. In addition, we have information about the gender, age, country and zip code of residence of the individuals. In Table 1, we present summary statistics for the demographic characteristics and the behavior of the individuals in our sample. The vast majority (93%) of individuals are men and the mean (median) age is 33 (31) years; these average characteristics are not very different from those reported for samples of individuals who invest in the stock market through online brokers (e.g., Barber and Odean, 2002; Glaser, 2003).¹¹ Each individual, on average, has wagered on 5 events during the same day, has wagered in the sportsbook on 35 different days, and has an average betting frequency of approximately once a week.

⁹Other studies (Pope and Peel, 1989; Levitt, 2004) suggest bookmakers may exploit individuals' biases by setting prices between the efficient ones and those that balance the book, but the bookmaker who provided our data has stated he does not use this strategy.

¹⁰The original dataset contains 400 individuals, but we drop 64 because they have placed fewer than 5 wagers.

¹¹In the Barber and Odean (2002) sample of U.S. investors, 86% of individuals are men and the mean (median) age is 49.6 (48) years. In the Glaser (2003) sample of German investors, 95% of individuals are men and the mean (median) age is 40.8 (39) years.

2.3 Lottery representation

In this section, we explain how we represent individual choices as lotteries, i.e., as probability distributions over monetary amounts (prizes). The simplest bet, which is a single bet, involves selecting an outcome in one event and has two possible prizes: if the selected outcome is realized, the prize equals the stake times the selection's return, otherwise the stake is lost. As in any field setting, it is not possible to know individuals' subjective beliefs, hence the probabilities individuals associate with these prizes. Since observed behavior can be explained by several combinations of risk preferences and subjective beliefs, this gives rise to the standard identification problem in studies that estimate risk preferences in the field. Virtually all such studies — whether they use data from horse-racetrack betting, insurance, options trading, or person-to-person lending — resolve this identification problem by assuming that beliefs are homogeneous and coincide with the rational, or more precisely, with the econometrician's approximation to the rational beliefs (see Barseghyan et al., 2015 for a review).¹² In the main analysis of our study, we make use of a key advantage of our market setting, which is that the true probabilities of the events can be readily approximated by the probabilities implied by their betting odds.¹³ We find that the market under study is quite efficient, hence the probabilities implied by the quoted prices are indeed close to the true probabilities of match outcomes (see Section 4.5 for details). Furthermore, the following features of our data suggest that individuals' subjective probabilities should not deviate significantly from these probabilities. First, we note that the betting platform provides access to online calculators that convert the quoted prices into their implied probabilities, so it is likely that individuals also observe the implied probabilities before placing their wagers. Second, we find that most individuals spread their wagers over various different leagues and teams: on average, an individual in our sample has placed wagers on 34 (182) different leagues (teams), while less than 20% (3%) of his wagers are placed on any specific league (team). Hence, it is unlikely that individuals have superior — real or perceived — information about any specific league/team, given that they only bet it so infrequently. In Section 4.5, we show that our results are robust to i) allowing for small but significant deviations of the subjective from the implied probabilities,

¹²This identification problem is, in fact, an issue in the wider theoretical and empirical literature in economics, where it is also commonly circumvented by assuming homogeneous rational beliefs; see Manski (2004) for a discussion.

¹³For comparison, in the insurance setting, it is typically assumed that each individual's subjective belief about his probability of having an accident equals the true probability, where the latter is estimated using a cross-sectional regression of claim rates on demographic characteristics (e.g., Barseghyan et al., 2013). In a TV game show in which the probability distribution of the prizes each contestant faces depends on the behavior of a counterparty, Post et al. (2008) propose a rule that explains 74% of the variance of the counterparty's behavior and assume that contestants' subjective beliefs are consistent with this rule.

and ii) to approximating the true probabilities with the win frequencies of past outcomes with similar prices.¹⁴

The more complex bets are *accumulators*, where an accumulator of type $k \geq 1$ involves wagers on k events. An accumulator has two possible prizes: if all wagers win, the prize equals the stake times the product of the selections' gross returns minus the stake, otherwise the stake is lost. The probability of the positive prize can be calculated as the product of the implied probabilities associated with each selection. For example, a bet that combines a wager at odds 2.0 with one at odds 2.5 is a type-2 accumulator, and has gross payoff equal to the stake times $2.0 \times 2.5 = 5.0$ if both wagers win and 0 otherwise; furthermore, if the implied win probabilities of the selections are 0.47 and 0.38, the probability of the high payoff is $0.47 \times 0.38 \approx 0.18$. Accumulators can also be combined into more complex bets, whose lottery representations can be constructed by combining the representations of the elemental accumulator bets. We create a lottery representation for each play session — defined as the set of all bets placed over a single day by each individual — by constructing all possible combinations of payoffs from all the bets placed during the session.¹⁵ After dropping 1,043 lotteries that are difficult to compute, we have a final sample of 11,490 lotteries chosen by 336 individuals.¹⁶ In Table 2, we report summary statistics for the characteristics of all the chosen lotteries. The median lottery contains 2 prizes, but more than 25% (5%) of lotteries contain 6 (61) prizes or more. The amount wagered ranges from €0.01 all the way to €5,500, and the maximum prize ranges from €0.01 to hundreds of thousands of euros. All lotteries have negative mean due to the bookmaker's commission that is embedded in the odds, while the summary statistics of the standard deviation and skewness demonstrate that individuals have chosen lotteries with a wide variety of characteristics, ranging from almost safe lotteries that yield a tiny payoff with probability close to 1, to highly skewed lotteries that yield a very high payoff with probability close to 0.

¹⁴The latter is the approach employed by Jullien and Salanie (2000) and Snowberg and Wolfers (2010) in studies that use data on parimutuel betting in the horse racetrack, a setting in which the market is quite inefficient hence implied probabilities deviate significantly from the true. Since prices are quite efficient in our fixed-odds betting market, and the implied probabilities are close to the win frequencies of past outcomes with similar prices, it is not surprising that our results are robust.

¹⁵Defining the play session as containing only individual bets would be too narrow since multiple bets are often placed simultaneously. At the other extreme, defining the play session as containing all bets that are not yet settled is very close to our definition of the play session, since most bets are placed on the day of the event.

¹⁶We drop days involving wagers on related events (e.g., the winner of a soccer match and the final score), because their lottery representation requires information about the events' joint distributions. As these days are relatively infrequent in our data and their occurrence is unlikely to be systematically related to our structural parameters, dropping them should have little effect on our results.

3 Reduced-form analysis

In this section, we conduct a preliminary panel-regression analysis to study the effect of previous bet outcomes on subsequent risk taking. We analyze the effect of individuals' cumulative past gains/losses on the extensive margin of risk taking by estimating a model of the decision to play in the sportsbook, and on the intensive margin of risk taking by estimating a model of the choice of lotteries conditional on playing.

The participation decision is modeled by the log-linear equation

$$\log(\text{Duration}_{nl}) = a_n + z'_{nl}b + \zeta_{nl},$$

where Duration_{nl} is the length (in calendar days) of the interval between the consecutive play days l and $l+1$ for individual n , z_{nl} is a vector of observables that include linear and quadratic terms of the cumulative gains and losses realized over a look-back period ending on (but not including) play day l , a_n denotes individual fixed effects, and b and ζ_{nl} are the regression coefficients and errors. This model essentially assumes that each explanatory variable accelerates/decelerates the time until the next play day by some constant, or equivalently, changes the scale but not the location of the baseline distribution of durations. We also note that by including cumulative gains and cumulative losses as separate explanatory variables in the model, we allow for the effect of past gains and losses to be different.

The model of the lottery-choice decision consists of a log-linear equation of the form

$$\log(Y_{nl}) = c_n + z'_{nl}d + \eta_{nl}$$

for each of four key lottery characteristics (mean, variance, skewness, and monetary stake), where Y_{nl} is a characteristic — more precisely, to ensure positivity, an affine transformation of a characteristic — of the lottery chosen by individual n on play day l , z_{nl} is as above, c_n denotes individual fixed effects, and d and η_{nl} are the regression coefficients and errors.

In Table 3, we present results from these two models, with previous gains and losses measured over a look-back period of 7 calendar days (with look-back periods of one month or one year yielding similar results). We observe that prior betting outcomes significantly affect both the frequency of playing and the characteristics of the chosen lotteries. In the participation model, cumulative gains and cumulative losses have a significantly negative effect, implying that individuals increase their play frequency following an increase in gains as well as in losses: For example, an increase of €100 in the cumulative gain (loss) reduces the length of the interval

between play days by 8% (3%). The quadratic terms of cumulative gains and losses have small but significantly positive coefficients, indicating that the effect of past outcomes on play frequency slowly tapers off. Similarly, in the lottery-choice model, we find that both past gains and past losses lead to higher risk taking, as individuals increase their stakes (by 14%/8% for €100 extra in gains/losses) and select lotteries that have higher variance (by 12%/9% for €100 extra in gains/losses). Interestingly, after prior losses but not after prior gains, individuals prefer lotteries with significantly higher skewness, possibly because these lotteries offer them a good opportunity to break even. As in the participation model, the small but significantly negative coefficients on the quadratic terms suggest that the effect of past outcomes on individuals' risk taking tapers off.

The results of the regression analysis indicate that individuals' choices are significantly affected by past outcomes, in a manner consistent with the house-money effect after gains and the break-even effect after losses. This behavior is at odds with the traditional framework of expected utility theory, and naturally lends itself to the prospect theory framework, where outcomes are evaluated relative to a reference point which can be history-dependent. In the next section, we estimate a structural model of prospect theory with history dependence, which can lead to a deeper understanding of these empirical relationships.

4 Structural analysis

In this section, we develop and estimate a model of individual bet choice that takes into account the play frequency and the lottery choice conditional on playing. This model-based framework enables us to link the empirical relationship between prior outcomes and subsequent risk taking to underlying economic primitives, namely individual-specific risk-preference parameters. First, we present the preferences that form the basis of our analysis. Second, we present our econometric implementation, which introduces heterogeneity across individuals and a random element in decisions that is necessary to explain the data with a(ny) theory of choice. Then, we present our results and robustness checks, and we provide intuition using illustrative examples.

4.1 Model of behavior

We model individuals' behavior as follows. On any given day, an individual has the opportunity to place a bet with some probability, where this probability depends on the individual's characteristics, such as how busy he is. If the individual has the opportunity to place a bet, he decides — based on his risk preferences — whether he wants to accept or reject this opportunity. For example, a more risk-averse individual may

reject betting opportunities more frequently, whereas a busier individual may have fewer such opportunities. If the individual rejects the opportunity to place a bet, he effectively chooses a safe lottery that pays 0 with certainty, while if he accepts it he chooses a risky lottery from the sportsbook.

In our analysis, we assume that individuals primarily care about the payoff risk profiles of lotteries rather than about the unique match characteristics such as, e.g., the specific teams involved. This assumption is motivated by the observation that individuals place a negligible proportion (on average, less than 3%) of their wagers on any specific team, and tend to combine wagers on various different sports, teams, and events, which indicates that they do not have strong favorites that affect their risk-taking decisions. As a result, even though different matches are available at different points in time, the set of available risky lotteries remains essentially unchanged, as at all points in time a multitude of prices, stakes, and bet types can be combined to produce the same universe of bets. Furthermore, we assume that individuals are unlikely to consider seriously lotteries that were *never* chosen by *any* individual, such as lotteries involving a large number of complex bet types.¹⁷ Hence, we construct the choice (or consideration) set of risky lotteries by randomly generating a large number of lotteries whose characteristics are drawn from the empirical distributions of the characteristics of the lotteries chosen by individuals in our data. In particular, we draw lottery characteristics using a procedure that resembles the procedure individuals employ when they place bets on the sportsbook, i.e., we select the number of bets to place during a play session, the type of each bet (e.g., single or combination bet), the amount of money to stake, and the odds associated with each wager.¹⁸ This choice set is then augmented with the chosen lottery and the safe lottery which represents the option not to place a bet.

Since wagering in the sportsbook involves a choice among lotteries, we need to specify individuals' preferences over lotteries. Motivated by our preliminary reduced-form analysis above, which shows that observed behavior can be naturally explained by a reference-dependent theory, we assume that preferences are represented by a variant of the cumulative prospect theory (CPT) proposed by Tversky and Kahneman (1992). In particular, at time t , an individual with preference parameter vector $\theta := (\alpha, \lambda, \gamma, \kappa, \delta)'$ evaluates lottery L with

¹⁷Even if individuals consider such lotteries, they likely have low utility so excluding them from the choice set should not affect our results; see Cohen and Einav (2007) and Barseghyan et al. (2013) for similar arguments for excluding alternatives from the choice set.

¹⁸See Section B in the appendix for details on the exact procedure and distributions used to generate the choice set. Also note that this procedure is essentially a reduced-form implementation of a structural model of choice set generation in which the complexity of a lottery decreases its likelihood of being included in the choice set; see Ben-Akiva et al. (1984) and Brownstone, Bunch and Train (2000) for analogous reduced-form approaches and Goeree (2008) for an analogous structural approach, in different contexts. Since the two approaches would yield similar choice sets, we have elected to use the reduced-form implementation to reduce computational complexity.

probabilities $\{p_i\}$ and payoffs $\{z_i\}$ with $z_{-m} \leq \dots \leq z_0 = 0 \leq \dots \leq z_M$ by computing the difference between his final wealth $z_i + W_t$ and a reference point W_t^{RP} , which depends on state variable s_t , and assigning to it utility

$$U(\theta, L, s_t) := \sum_i w_i v(z_i + W_t - W_t^{\text{RP}}). \quad (1)$$

In the above equation, $v(\cdot; \alpha, \lambda)$ is the value function,

$$w_i := \begin{cases} w(p_i + \dots + p_M) - w(p_{i+1} + \dots + p_M) & \text{for } 0 \leq i \leq M \\ w(p_{-m} + \dots + p_i) - w(p_{-m} + \dots + p_{i-1}) & \text{for } -m \leq i < 0 \end{cases} \quad (2)$$

is the decision weight for payoff z_i , with $w(\cdot; \gamma, \kappa)$ the probability weighting function, and $W_t - W_t^{\text{RP}}$ is the difference between current and reference wealth at time t and depends on s_t and δ .¹⁹

For the value function, we use the form proposed by Tversky and Kahneman (1992):

$$v(x; \alpha, \lambda) := \begin{cases} x^\alpha & \text{for } x \geq 0 \\ -\lambda(-x)^\alpha & \text{for } x < 0 \end{cases}. \quad (3)$$

The curvature parameter $\alpha \in (0, 1]$ measures diminishing sensitivity to deviations from the reference point, which implies that the value function is concave over gains and convex over losses. The lower α is, the higher the curvature of the value function, i.e., the faster the sensitivity to deviations from the reference point diminishes. The loss aversion parameter $\lambda > 0$ measures the relative sensitivity to gains versus losses. Values $\lambda > 1$ imply a higher sensitivity to losses than to gains (loss aversion), while values $0 < \lambda < 1$ imply a higher affinity to gains (gain seekingness).²⁰

For the probability weighting function, we use the two-parameter function suggested by Lattimore, Baker and Witte (1992), which is flexible enough to separately capture the curvature and the elevation of the weighting function, and has been shown to account well for individual heterogeneity. That is, we use

$$w(p; \gamma, \kappa) := \frac{\kappa p^\gamma}{\kappa p^\gamma + (1-p)^\gamma}, \quad (4)$$

where parameter $\gamma > 0$ measures the curvature and parameter $\kappa > 0$ measures the elevation of the probability weighting function. Intuitively, γ captures diminishing sensitivity to changes in cumulative probability from

¹⁹Note that, in this subsection, we simplify notation by omitting the subscript n denoting the individual.

²⁰In the original Tversky and Kahneman (1992) CPT formulation, λ is by assumption constrained to be greater than one (hence termed the loss aversion parameter) to reflect the observation that “losses loom larger than gains” (e.g., the endowment effect and the status quo bias). We do not place this restriction on λ , because recent experimental evidence (e.g., Abdellaoui, Bleichrodt and Paraschiv, 2007; von Gaudecker, van Soest and Wengstrom, 2011) has found that a small but significant proportion of individuals exhibit gain seekingness, i.e., $\lambda < 1$. As we discuss in Section 4.3, we estimate that only a small proportion (about 10%) of individuals have $\lambda < 1$; re-estimating the model with the constraint $\lambda \geq 1$ has no significant effect on our results.

the natural boundaries of certainty and impossibility, and κ captures the sensitivity to extreme outcomes. In Figure 2, we show the effects of γ and κ on the probability weighting function and its derivative. Setting $\kappa = 1$ and considering the effect of γ , we see that for $\gamma < 1$ the probability weighting function has an inverse-S shape which corresponds to an overweighting of outcomes in the tails of the distribution relative to intermediate outcomes; the lower the value of γ , the more pronounced is this shape. For $\gamma > 1$ (not shown), the function becomes S-shaped and underweights the tails. Setting $\gamma = 1$ and considering the effect of κ , we see that for $\kappa > 1$ the probability weighting function is globally concave which corresponds to an overweighting of the right (left) tail of the distribution of positive (negative) outcomes; as κ gets higher, the concavity of the function increases. When $\kappa < 1$ (not shown), the function becomes globally convex.²¹

The original CPT pertains to static choice settings and assumes that the reference point that separates losses from gains is equal to the status quo, i.e., the individual's current wealth ($W^{\text{RP}} := W$). Motivated by experimental evidence that prior outcomes affect individuals' subsequent risk taking (Thaler and Johnson, 1990), a number of studies (e.g., Barberis, Huang and Santos, 2001; Post et al., 2008; Barberis and Xiong, 2009) extend prospect theory in a dynamic setting by assuming that prior outcomes affect the reference point hence the framing of subsequent payoffs as gains and losses relative to it. For example, Barberis and Xiong (2009) assume that the reference point does not update after gains/losses but rather completely sticks to its previous level, while Post et al. (2008) allow the reference point to be a linear function of the two most recent past expectations and the contemporaneous expectation of future outcomes. Following an approach that is similar but more parsimonious, and utilizing the standard exponential-discounting modeling assumption, we allow the reference point on date t to depend on previous outcomes according to

$$W_t^{\text{RP}} := W_t - \sum_{k=1}^{K_t} \delta^k PNL_{t,k}, \quad (5)$$

where $PNL_{t,k}$ is the gain/loss realized on the k -th day (counting backward from date t) on which a gain or loss was realized, K_t is the number of such days from the beginning up to (but excluding) date t , and $0 \leq \delta \leq 1$ is a discount factor. That is, we assume that individuals attach exponentially lower weights to more distant outcomes, with parameter δ the individual-specific memory decay parameter which controls how far back

²¹To reduce the number of model parameters and facilitate identification, we deviate from the original Tversky and Kahneman (1992) CPT formulation in which the value and probability weighting functions have different parameters for gains and losses. Indeed, Tversky and Kahneman (1992) estimate these parameters to be very similar for gains and losses, so many studies that estimate or use CPT preferences make similar simplifying assumptions.

individuals recall past gains and losses. When δ is low, the reference point is close to contemporaneous wealth, as the original CPT suggests: previous gains and losses are quickly absorbed and do not affect behavior for long. When δ is high, the reference point does not quickly update to incorporate previous outcomes but rather sticks in between previous and current wealth levels. In this case, the individual has a long memory, which makes him perceive losses that follow (larger) prior gains as a reduction in the prior gain rather than as a loss, and gains that follow (larger) prior losses as a reduction in the prior loss rather than as a gain.²²

The relatively large number of choices we observe for each individual, coupled with the wide variety of lotteries from which individuals make their choices, facilitates the identification of the model parameters. First, we can identify between risk aversion and betting opportunities because both affect betting frequency but only the former affects lottery choice conditional on having the opportunity to place a bet. Second, it is possible to identify between risk aversion, i.e., value-function curvature, and loss aversion due to the existence of both small- and large-stakes lotteries: A risk-averse individual who is averse to small-stakes lotteries would be extremely averse to large-stakes lotteries, as risk aversion affects the global concavity of the value function, while a loss-averse individual who is averse to small-stakes lotteries would not be as averse to large-stakes lotteries, as loss aversion increases primarily the local concavity of the value function at the reference point (see Rabin, 2000). Third, we can identify between risk aversion and probability weighting due to the existence of lotteries with different levels of variance and skewness: An individual who overweights small probabilities would prefer lotteries with high positive skewness, while a risk-loving individual would prefer lotteries with high variance. Fourth, the existence of lotteries with a wide range of distributions over gains and losses helps us identify between the curvature and elevation of the probability weighting function. As illustrated in Figure 2, an individual with a curved probability weighting function overweights outcomes at both tails of the distribution (both for gains and losses) at the expense of outcomes in the middle, while an individual with an elevated probability weighting function overweights outcomes at one tail of the distribution (the left for losses and right for gains) at the expense, primarily, of outcomes at the other tail. Finally, to identify the dependence of the reference point on past gains and losses, we exploit our parsimonious assumption that the dependence takes the familiar exponential-discounting form as well as the relatively long time dimension of our panel, which enable us to uncover systematic time variation in individuals' choices.

²²Though past gains/losses affect the reference point symmetrically in our model, this setup is flexible enough to accommodate an asymmetric response of risk taking to past gains/losses, hence behaviors like the house-money and break-even effects (see Section 4.6).

4.2 Econometric implementation

Based on this model of behavior, here we present our structural econometric model, which introduces a stochastic element in decisions and individual heterogeneity. In particular, we use a random-utility discrete-choice model (Marschak, 1960; McFadden, 1974) that incorporates two key ingredients: first, model parameters vary across individuals, and second, they are drawn from a population distribution.

In the random-utility model, the utility V_{njt} that individual n gets from choosing lottery L_j on date t is decomposed into a deterministic and a random component as

$$V_{njt} := U(\theta_n, L_j, s_{nt}) + \varepsilon_{njt},$$

where $\theta_n := (\alpha_n, \lambda_n, \gamma_n, \kappa_n, \delta_n)'$ is the preference parameter vector of individual n , $s_{nt} := \{PNL_{nt,k}\}$ is the sequence of gains/losses realized by individual n until date t , and ε_{njt} is the random component. The random component can account for mistakes the individual makes in evaluating or choosing an alternative, or can simply be interpreted as an econometric error necessary to reconcile the data with a theory of choice. As is standard in the literature on discrete choice, we assume the error terms are i.i.d. and follow the double exponential distribution with location parameter normalized to 0 (without loss of generality) and inverse scale parameter $\tau_n > 0$ for individual n ; this is a symmetric distribution with shape similar to the normal but with heavier tails, yielding more robust analysis. As a result, the probability that an individual chooses lottery L_j on date t out of the set of alternative lotteries C is

$$\begin{aligned} p(y_{nt} = L_j | \tau_n, \theta_n, s_{nt}) &= p\left(U(\theta_n, L_j, s_{nt}) + \varepsilon_{njt} > \max_{i \neq j \in C} \{U(\theta_n, L_i, s_{nt}) + \varepsilon_{nit}\}\right) \\ &= \frac{\exp(\tau_n U(\theta_n, L_j, s_{nt}))}{\sum_{i \in C} \exp(\tau_n U(\theta_n, L_i, s_{nt}))}. \end{aligned} \quad (6)$$

To reduce the computational complexity due to the choice set having a large number of similar alternatives (so a large number of costly utility evaluations in the denominator in Equation 6), we first reduce the number of alternatives by grouping them into 100 clusters according to a measure of distribution similarity and keeping the most representative lottery in each cluster, and then add the chosen and the safe lottery. We note that increasing, e.g., to 200, the number of clusters has a very small effect on our results; also see Section B of the appendix for details on the clustering algorithm.

We model heterogeneity across individuals in all parameters: the preference parameters θ_n , the inverse

scale τ_n of the random choice errors, and the probability π_n of having the opportunity to place a bet on any given day. For convenience, we denote the model parameters collectively by $\vartheta_n := (\pi_n, \tau_n, \theta'_n)'$. Recognizing that individuals form a sample from a population, we model these individual-specific parameters as draws from a distribution. In particular, since the elements of ϑ_n are bounded, we define transformed parameters $\tilde{\vartheta}_n := g(\vartheta_n)$, where $g(\cdot)$ maps — through Johnson transformations — elements of ϑ_n to $(-\infty, +\infty)$, and assume that the $\tilde{\vartheta}_n$ are i.i.d. draws from a multivariate normal with population mean and variance $\mu_{\tilde{\vartheta}}$ and $\Sigma_{\tilde{\vartheta}}$. That is, $\forall \vartheta_n$

$$g(\vartheta_n) | \mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}} \sim \mathcal{N}(\mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}}). \quad (7)$$

As a result, the population mean and variance — a measure of heterogeneity — are introduced as model parameters, for which we can draw inference using the data.²³

As already mentioned, to study behavior we utilize information in individuals' observed choices, i.e., both the wagers they place and the frequency with which they place wagers. In particular, we let x_{nt} be a dummy indicating whether individual n played on date t and y_{nt} the lottery choice we observe for individual n on date t . If individual n has the opportunity to bet on date t , he either i) chooses the safe lottery L_0 that pays 0 with certainty, hence we observe no bet on date t , i.e., $x_{nt} = 0$, or ii) he chooses lottery $y_{nt} \neq L_0$, hence we observe a bet on date t , i.e., $x_{nt} = 1$. If he does not have the opportunity to bet, we observe no bet on date t , i.e., $x_{nt} = 0$. Thus, the likelihood of observing (x_{nt}, y_{nt}) given the individual's parameters and state is

$$p(x_{nt}, y_{nt} | \pi_n, \tau_n, \theta_n, s_{nt}) = x_{nt} \pi_n p(y_{nt} | \tau_n, \theta_n, s_{nt}) + (1 - x_{nt}) [\pi_n p(y_{nt} = L_0 | \tau_n, \theta_n, s_{nt}) + (1 - \pi_n)], \quad (8)$$

where $p(y_{nt} | \tau_n, \theta_n, s_{nt})$ is given by Equation 6 above.

For tractability, we estimate the model using Bayesian techniques. For this reason, we augment the model with priors for the population parameters $\mu_{\tilde{\vartheta}}$ and $\Sigma_{\tilde{\vartheta}}$. The joint prior density of all model parameters is

$$p(\{\vartheta_n\}, \mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}}) = \left\{ \prod_n p(\vartheta_n | \mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}}) \right\} \times p(\mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}}), \quad (9)$$

where $p(\vartheta_n | \mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}})$ follows from Equation 7 using the normal density and the Jacobian method, and for $p(\mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}})$ we adopt the standard independent Normal-inverse-Wishart conjugate prior, i.e.,

$$\mu_{\tilde{\vartheta}} \sim \mathcal{N}(\underline{\kappa}, \underline{K}) \quad \Sigma_{\tilde{\vartheta}}^{-1} \sim \mathcal{W}(\underline{\lambda}, \underline{\Lambda}^{-1}). \quad (10)$$

²³Von Gaudecker, van Soest and Wengstrom (2011) and Barseghyan et al. (2013) find that, both in the lab and in the field, heterogeneity in individual risk preferences is almost entirely unobserved. As a result, to aid tractability and identification, we do not introduce an explicit dependence of the model parameters on the individual characteristics we observe (age, gender, and zip code).

Our baseline priors are $\underline{\kappa} = \mathbf{0}$, $\underline{K} = 100I$, $\underline{\lambda} = \text{rank}(\Sigma_{\hat{\gamma}})$, and $\underline{\Lambda} = I$. Setting \underline{K} large and $\underline{\lambda}$ small makes the priors weak and lets the data determine the posteriors. Though this renders the choice of $\underline{\kappa}$ immaterial, posteriors could still be affected by $\underline{\Lambda}$, so we perform a prior sensitivity analysis which shows they are robust.

To estimate the model using Bayesian techniques, we need to obtain information about the joint posterior distribution of the model parameters conditional on the data. The joint posterior is proportional to the likelihood (Equation 8) times the joint prior (Equation 9) but cannot be calculated analytically, so we make draws from it using a Markov Chain Monte Carlo algorithm with Metropolis-within-Gibbs sampling (Chib and Greenberg, 1995). We delegate a detailed description of our algorithm to Section A in the appendix. Next, we present our results based on the posterior distribution of the model parameters.

4.3 Estimation results

In Panel A of Table 4, we present the posterior estimates for the population means of the model parameters, and in Table 5 we also show the medians of the estimated population distributions.

Our population estimates for the CPT parameters imply that, on average, the concavity (convexity) of the value function in the region of gains (losses) is mild (the mean/median of α is 0.86/0.87), and there is a modest degree of loss aversion (the mean/median of λ is 1.48/1.37). Furthermore, our population estimates for the curvature and the elevation of the probability weighting function (the mean/median of γ is 0.88/0.91 and of κ is 1.41/1.17) imply that, on average, the probability weighting function is concave for most probabilities and has a slight inverse-S shape, resulting in significant overweighting of extreme positive and extreme negative outcomes. In Figure 4, we plot the value function and the probability weighting function corresponding to various parameter combinations drawn from the posterior population distributions of parameters α , λ , γ , and κ . Using a set of hypothetical lotteries, Tversky and Kahneman (1992) estimate that, on average, the value function has similar curvature (median α is 0.88) to the one we estimate and a more pronounced kink at the reference point, corresponding to stronger loss aversion (median λ is 2.25), and that the average probability weighting function has a more pronounced inverse-S shape hence additionally overweights outcomes near the reference point. Since the original experiment of Tversky and Kahneman (1992), various experimental studies have reported a wide range of estimates for the CPT parameters: According to a meta-analysis study by Booij, van Praag and van de Kuilen (2010), the average curvature

parameter α of the value function has been estimated in the range [0.2, 0.9] and the average loss aversion parameter λ in the range [1.07, 3.2].²⁴ For the probability weighting function, different studies have estimated different functional forms, whose parameters usually indicate an inverse-S shape, while a few studies do not find support for the changing concavity of the probability weighting function (e.g., Abdellaoui, Vossman and Weber, 2005). Thus, though some of our estimates deviate from the widely used estimates found in the Tversky and Kahneman (1992) experiment, they are well within the range of estimates found in the experimental literature.²⁵ We conclude then that our findings complement the general experimental finding that individuals are mildly risk averse (loving) over gains (losses), moderately loss averse, and overweight the probabilities of extreme outcomes, and they extend support for this finding to a real market setting.

Regarding the memory decay parameter δ , we estimate its population mean/median (0.60/0.66) to be significantly different from zero, which indicates that most individuals do not quickly absorb past profits and the reference point they use to separate gains from losses tends to stick to previous levels. For example, $\delta = 0.66$ implies that an individual's reference point absorbs 34%/56%/87% of the gain/loss realized 1/2/5 play days earlier. That is, after a gain of, e.g., €10 the reference point is $(100\% - 34\%) \times €10 = €6.6$ below the actual wealth level, so a subsequent loss smaller than €6.6 is evaluated as a gain. To our knowledge, there are no estimates in the literature for this memory decay parameter. Given the lack of previous estimates of a structural relationship, dynamic models in behavioral finance have so far incorporated this framing effect by arbitrarily assuming that individuals have a perfectly sticky reference point, i.e., $\delta = 1$. As we show in Section 5 on the disposition effect, the assumption that $\delta = 1$ may lead to qualitatively different conclusions than an assumption informed by actual estimates, and furthermore, it cannot be thought of as corresponding to one extreme in behavior (with $\delta = 0$ corresponding to the other extreme) because δ may affect behavior non-monotonically.

In Panel B of Table 4, we present the posterior estimates for the population variances of the model parameters, in Table 5 we present percentiles of the estimated distributions of the model parameters, and

²⁴In particular, out of the experiments analyzed by Booij, van Praag and van de Kuilen (2010), 8 studies estimate an average α in the range [0.2, 0.7] and 4 estimate it in [0.8, 0.9], while 3 studies estimate an average λ below 1.5 and 3 estimate it above 2.5 (note that half the studies do not estimate λ because they do not consider mixed lotteries).

²⁵As Fehr-Duda and Epper (2012) point out, it is puzzling that the literature has viewed with suspicion parameter estimates that deviate from those of Tversky and Kahneman (1992), despite the fact that Tversky and Kahneman (1992) i) report median estimates without reporting standard errors on their estimates or any measure of heterogeneity across individuals, and ii) base their estimates on the decisions of a small sample of 25 students, who chose between hypothetical lotteries that involved only two outcomes and small stakes. While it is understandably convenient to have one unambiguous set of parameter estimates to use in applications, the reality is that there is wide heterogeneity in estimates both across individuals and across studies.

in Figure 3 we plot the densities of these estimated distributions. We find substantial heterogeneity across individuals in all preference parameters, which is mostly ignored by applied models in finance. In particular, our estimates imply that the 25th/75th percentiles of the posterior distribution of the value function parameters are 0.83/0.90 for α and 1.18/1.70 for λ , and the corresponding percentiles for the probability weighting function parameters are 0.85/0.94 for γ and 1.00/1.59 for κ . In Figure 4, we present plots of the value and probability weighting functions corresponding to these percentiles. Regarding the memory decay parameter δ , its posterior distribution reveals that there is a continuum of individuals who are affected to different extents by previous outcomes, but with a marked concentration close to the two natural extremes of no memory and perfect memory. In particular, 15% of the individuals in our sample have very short memory ($\delta < 0.1$), 22% of the individuals have very long memory ($\delta > 0.9$), and the remaining 63% of the individuals have memory that decays at a wide range of rates as evidenced by the distance between the 25th and 75th percentiles (at 0.34 and 0.87, respectively) of the estimated distribution of δ . We note that the heterogeneity we estimate in all model parameters is not driven by the choice of our priors, since in a prior sensitivity analysis we find that the priors have a very small effect on our posteriors. In Figure 3, we also show the results of the prior sensitivity analysis with respect to the scale $\underline{\Lambda}$ of the prior distribution of $\Sigma_{\bar{\theta}}$, which is the only informative prior that we use in our analysis, hence is the likeliest to have an effect on the posteriors; we find that the estimated population distributions for all model parameters are not significantly affected when we replace our baseline prior ($\underline{\Lambda} = I$) with either a low-variance prior ($\underline{\Lambda} = 0.1I$) or a high-variance prior ($\underline{\Lambda} = 10I$).

4.4 Model fit

As discussed above, a random utility component is necessary to reconcile individuals' repeated choices with any deterministic theory of choice which predicts a unique optimal choice for a given set of preference parameters. This random component represents omitted factors that may affect the utility of each lottery and make an individual more or less likely to choose an alternative relative to what his risk preferences would predict (e.g., utility from gambling, preferences over unique match characteristics). As a result, the magnitude of this component provides a sense of the extent to which unobserved factors are crucial in explaining the observed choices, and therefore the extent to which our model is a "good" description of individual behavior. In Panel A of Table 4, we see that the population mean of the inverse scale of the random choice parameter τ is 1.36.

Having normalized to one the utility difference between lotteries that yield a gain of €0 and €1 with certainty, this estimate implies that, e.g., an individual choosing between two lotteries whose utility differs by 1 (by 10), chooses the less-preferred one with probability 0.20 ($1.24 \cdot 10^{-6}$). Hence, the random component needed to fit our model to the observed choices is relatively small, and comparable to that estimated in experimental studies that structurally estimate preferences using a similar number of observations per individual (e.g., von Gaudecker, van Soest and Wengstrom, 2011). Nonetheless, this small random utility component plays an important role in explaining both the variation in individuals' risky choices over time as well as why individuals sometimes participate in the sportsbook — i.e., choose a risky lottery over the safe alternative — and sometimes do not. For example, we find that for a sizable proportion of individuals (about two thirds) the certainty equivalent of the chosen lottery is very slightly negative. This implies that, deterministically, these individuals are almost indifferent between the safe and the chosen lottery, and a very small random utility component (e.g., from utility of gambling) is sufficient to explain their participation decision.²⁶ Furthermore, in Figure 5 we plot the histogram of the mean — across all choices for each individual — ratio ranking of the chosen lotteries. We see that the chosen lotteries are rarely the top-ranked alternative according to the deterministic component of utility so a random component is necessary to explain the variation in observed choices. But this random component is quite small, as we have stated above, and indeed for the majority of individuals, the chosen lotteries are among the top 10% of the most-preferred alternatives, and for about half the individuals they are among the top 5%.

We also compare the fit of our model with that of several alternative models: i) one in which lotteries are chosen randomly; ii) one in which individuals are risk neutral, iii) one in which preferences are homogeneous, and iv) one in which preferences are heterogeneous but do not depend on previous outcomes, i.e., all individuals have no memory. In Table 6, we see that the average log likelihood from our model is significantly higher than that of the alternative models, suggesting that both preference heterogeneity and history dependence play a significant role in explaining the observed behavior. The Deviance Information Criterion, which penalizes model complexity, is lowest for our model hence it also selects our model which incorporates both heterogeneity and history dependence.

²⁶To be specific, the maximum (across individuals) difference between the safe lottery and the individual's most-preferred risky lottery is just €0.05 in certainty equivalent terms. Notably, the random utility component that is sufficient to explain individuals' choice between risky lotteries and the safe lottery is smaller than that needed to explain the variation in individuals' risky choices over time.

4.5 Subjective beliefs

As discussed in Section 2.3, differences in behavior can be explained by several combinations of preferences and beliefs, which poses an identification problem in all studies that estimate preferences in the field. In our main analysis, we resolve this problem by using the standard assumption that beliefs are rational, which we approximate with the probabilities implied by the quoted odds. In this section, we first show that the market under study is quite efficient, hence the probabilities implied by the prices are indeed close to the true probabilities of match outcomes. Then, we show that our results are robust to i) approximating the true probabilities of bet outcomes with the win frequencies of past outcomes with similar prices rather than with the implied probabilities, and ii) allowing for small but significant deviations of the subjective from the implied probabilities.

To examine the efficiency of the quoted prices, we obtain historical data on the odds and results of all soccer matches that were offered by the bookmaker under study during our sample period.²⁷ We divide all possible outcomes of these matches into 100 percentile odds groups, and in Figure 6 we plot for each of these groups the mean realized rate of return in excess of the mean rate of return across all odds groups, along with the corresponding 95% bootstrap confidence intervals. The plot does not reveal large and/or significant inefficiencies in the quoted prices. Though there is a negative slope indicating a mild favorite-longshot (FL) inefficiency, i.e., the return from betting on short odds is higher than that from betting on long odds, this is not very pronounced. In particular, there are a number of longshot percentiles with the same estimated excess return as the favorite percentiles, while the 95% confidence intervals show that the excess return for most percentiles is not significantly different from zero. Indeed, the FL inefficiency that we find is much less pronounced than that documented in studies that use data on betting in parimutuel horse races (see, for comparison, Figure 6, in which we superimpose the corresponding plot for the Snowberg and Wolfers, 2010 data on parimutuel betting in the horse racetrack).²⁸ Nonetheless, to examine whether the mild FL inefficiency that is present in our market affects our results, we repeat our analysis replacing the implied probabilities with the win frequencies of past outcomes with similar prices. In Figure 7a, we see that our

²⁷We obtained the data from the following sources: i) football-data.co.uk, ii) matchstatistics.com, and iii) betfair.com. We restrict this analysis to the soccer wagering market because historical data are significantly more readily available for soccer than for other sports. Furthermore, this is the most active market segment with the highest transaction volume in the sportsbook we study.

²⁸We note that, while the literature has consistently found large pricing inefficiencies in parimutuel horse-racetrack betting markets, the evidence on the efficiency of fixed-odds sports betting markets is mixed. For example, Cain, Law and Peel (2000) and Kuypers (2000) find inefficiencies sufficient to allow for positive returns, while Pope and Peel (1989) and Woodland and Woodland (1994) find that the setting of odds is quite efficient, which is what we also find in our market.

results remain largely unchanged under this alternative assumption.

To examine the sensitivity of our results to our identification assumption of rational beliefs, we repeat our analysis allowing observed choices to be (at least partially) motivated by subjective beliefs. In particular, letting $1-p$ be the implied win probability for an outcome selected by an individual, we allow the “perceived” win probability of this outcome to be equal to $1-\zeta p$ for $\zeta < 1$. In Figure 7b, we plot the posterior distributions of the preference parameters for two different values of ζ ; small (large) deviations yield an average, across all lotteries in our sample, increase of 1.4% (3.3%) in the perceived probability of winning the maximum prize, and an average increase of 2.2% (4.5%) in the lottery’s expected value. We see that deviations of the subjective from the implied probabilities do not significantly affect our estimates. The reason is that, even if an individual places a wager on an event because he thinks he has superior information (e.g., he thinks that Barcelona is better than the market seems to think), he still needs to choose among a multitude of wagers involving this event (e.g., that Barcelona will win or that it will win by a large margin), where each wager involves a different level of risk. Furthermore, individuals typically place multiple wagers on the same day, hence they also need to decide if and how they want to combine these wagers under different bet types involving different levels of risk. Hence, even though our assumption of rational beliefs may not be exactly correct for all selections of all individuals, the chosen day lotteries are still informative about individuals’ risk preferences.

4.6 Intuition — Examples

Here, we present intuition about the effects that the CPT preference parameters and different levels of past profits have on individuals’ subsequent risk taking. We focus on three interesting cases of individuals from our sample for whom we have estimated different value-function curvature, loss aversion, and memory decay parameters. We analyze their optimal behavior in two hypothetical states in which they face the option of accepting or rejecting the same binary lottery with one negative and one positive payoff; in one state, individuals have no previous profits (neutral state), and in the other state they have previously realized a gain (gain state), which can be either “small” or “large” relative to the payoffs of the lottery.

The first individual we consider (Individual *A*) is estimated to have a mildly curved value function ($\alpha = 0.88$), to be moderately loss averse ($\lambda = 1.41$), and to use a reference point that perfectly sticks to previous wealth levels ($\delta \approx 1$). The second individual (Individual *B*) is estimated to have value function parameters

similar to Individual *A* but to use a reference point that partially updates to incorporate previous outcomes ($\delta = 0.50$). The third individual (Individual *C*) is estimated to have a perfectly sticky reference point, similar to Individual *A*, but a highly curved value function ($\alpha = 0.50$) and a low degree of loss aversion ($\lambda = 1.12$). To facilitate the exposition, we select these three individuals such that their estimated probability weighting parameters γ and κ are close to 1, i.e., none of them distorts event probabilities. The binary lottery we consider has payoffs -7 or 10 , with equal probability, chosen such that all individuals we consider are indifferent between choosing the lottery or not in the neutral state, which facilitates our comparison of risk taking across the different states.

In Panels (a), (c), and (d) of Figure 8, we plot for each individual his value function and the certainty equivalent of the lottery in the neutral state and in the gain state after a large prior gain of 8. In the neutral state, the reference point of all individuals is equal to their current wealth and they face a lottery whose payoffs straddle the reference point; in this state, loss aversion plays a dominant role in the individuals' decision, because the kink in the value function at the reference point is more pronounced locally than its curvature. As already mentioned, all individuals are by design indifferent between choosing the lottery or not in the neutral state, so we focus on comparing their behavior in the gain state. Individuals *A* and *C* do not integrate the gain to their reference point which remains at their original wealth (before this gain), hence they perceive the lottery as having possible outcomes $8 - 7 = 1$ or $8 + 10 = 18$ relative to this reference point. Since both lottery payoffs are perceived as relative gains, loss aversion is irrelevant and the concavity of the value function over gains becomes the dominant effect in the decision. If this concavity is mild enough relative to loss aversion, which is the case for Individual *A*, the lottery appears more appealing than in the neutral state; if it is strong enough, which is the case for Individual *C*, the lottery appears less attractive. As a result, Individual *A* exhibits increased risk taking after gains — consistent with the house-money effect — while Individual *C* exhibits the opposite behavior. Contrary to Individual *A*, Individual *B* integrates half of the gain into his reference point, which therefore updates halfway between his wealth before and after this gain, and therefore he perceives the lottery as having possible outcomes $4 - 7 = -3$ or $4 + 10 = 14$ relative to this reference point. Since the perceived lottery payoffs straddle the reference point, the effect of loss aversion is compounded by that of the value function curvature, hence for Individual *B* the lottery appears less appealing than in the neutral state.

In Panel (b) of Figure 8, we plot for Individual *A* his value function and the certainty equivalent for the lottery in the neutral state and after a small prior gain of 4. While Individual *A* accepts the lottery after a

large gain, he rejects it after a small gain. Comparing the behavior of Individual A after a small gain with the behavior of Individual B after a large gain, we see that a smaller prior gain has the same effect on individual behavior as a less sticky reference point, as in both cases the (perceived) lottery payoffs in the gain state straddle the reference point in a way that the interplay of loss aversion and risk aversion makes the lottery less appealing than in the neutral state.

A similar line of reasoning can be applied to analyze individual risk-taking behavior after losses. For example, after a large loss, Individual A will exhibit increased risk taking — consistent with the break-even effect — as the convexity of the value function over losses becomes the dominant effect on his decision, while the opposite effect would be observed after a small loss and/or for individuals with different preference parameters.

These examples demonstrate that in the presence of heterogeneity in individuals' risk preferences and in the level of prior profits, estimating an average reduced-form empirical relationship between past profits and subsequent risk taking can yield misleading results. This highlights the importance of structural estimation of heterogeneous risk preferences, which can help us better understand and predict behavior. In the following section, we apply our estimates to study the dynamic trading behavior of individuals in the stock market.

5 An application to the disposition effect

Applied models in finance have used prospect theory to explain asset prices, portfolio choice, trading decisions over time, and other behaviors (see Barberis, 2013 for a review). The results of these studies are largely based on i) a representative-agent assumption, ii) the median CPT parameter estimates from the Tversky and Kahneman (1992) experiment, and — in dynamic settings — iii) the assumption that individuals have either no memory or perfect memory. In this section, we use the individual-level preference parameters we estimate in Section 4 to study individuals' dynamic trading behavior, with a particular focus on the disposition effect, i.e., the tendency of individuals to sell (hold) stocks whose value has increased (decreased) since purchase. Our goal is to understand the implications of replacing the usual assumptions with our estimates which are derived from real data and incorporate heterogeneity and history dependence. In particular, we study how the preference heterogeneity we estimate translates to heterogeneity in exhibiting the disposition effect, how the history dependence in the estimated preferences affects the disposition effect, and more generally, whether prospect theory can explain the prevalence of this behavior.

The disposition effect (henceforth DE) is a very robust empirical finding regarding trading behavior.²⁹ It has long been argued informally (e.g., Weber and Camerer, 1998) that prospect theory can explain the DE, as follows. If a prospect-theory investor holds a stock that has gained (lost) value since purchase, continuing to hold it involves a gamble over gains (losses) relative to his original reference point, i.e., the purchase price. And since he is risk-averse (risk-seeking) over gains (losses), he prefers to sell (hold on to) the stock, therefore he exhibits the DE. However, this argument ignores the effect of loss aversion on investors' decisions and how it interacts with the asset's risk characteristics. This is important because, in prospect theory, loss aversion rather than the curvature of the value function is the primary determinant of behavior toward (relatively) small-stakes gambles. Barberis and Xiong (2009) use a simple model to formally show that — contrary to the conventional wisdom — CPT investors with preference parameters as estimated by Tversky and Kahneman (1992) and a perfectly sticky reference point would optimally increase (decrease) their holdings of a stock after it appreciates (depreciates), which is the opposite behavior to the DE.³⁰ The intuition behind this is the following. When the investor first decides whether to buy a stock or not, he faces a lottery whose payoffs straddle the reference point, hence loss aversion “bites” and can dominate the effect of value-function curvature on the decision. As a result, the investor only buys a stock with a sufficiently appealing risk profile, which for stocks plausibly translates to larger potential gains than potential losses. If the investor buys such a stock and it subsequently experiences a gain, then if his reference point is perfectly sticky he subsequently faces a lottery whose payoffs are all relative gains, so loss aversion is irrelevant and the mild concavity of the value function over gains becomes the dominant effect. This implies that the prospect of holding on to the stock is more appealing than that of buying it in the first place, hence the investor does not exhibit the DE. For an illustration of this intuition, see Figure 8a, which depicts the choice to take risk after a “large” gain.

Model setup Following Barberis and Xiong (2009) (henceforth BX), we consider a discrete-time binomial model with dates $t \in \{0, 1, \dots, T\}$, where 0 is the beginning and T the end of the year. At each date $t \in \{0, \dots, T-1\}$, there is a risk-free asset with (gross) return R_f and a risky asset with (gross) return $R_t \in \{R_u, R_d\}$ such that its excess return has annualized mean μ and standard deviation σ , with

²⁹See Odean (1998), Grinblatt and Keloharju (2001), and Feng and Seasholes (2005) for evidence using individual trading data for U.S., Finnish, and Chinese investors, respectively; Frazzini (2006) for evidence on U.S. mutual fund managers; and Weber and Camerer (1998) for evidence using experimental data.

³⁰See also Hens and Vlcek (2011) for similar results using a slightly different model.

$0 < R_d < R_f < R_u$ (to preclude arbitrage). To maintain comparability with the results in BX, we set the risk-free return R_f to 1, the mean μ in the range 3% to 13%, and the standard deviation σ to 30% — which together correspond to a range of reasonable values for the equity premium and Sharpe ratio — and $\Pr(R_t = R_u)$ to 0.5.³¹ Furthermore, we set $T = 2$, i.e., there are three dates and two trading periods in the year, as this corresponds well to the empirically observed trading frequency for individual investors (Barber and Odean, 2000; Alvarez, Guiso and Lippi, 2012), and is useful in terms of tractability and simplicity of exposition.³²

At $t = 0$, the unit price of the risky asset is (an arbitrary) $P_0 > 0$ and the investor has initial wealth (an arbitrary) $W_0 > 0$.³³ At dates $t \in \{0, 1\}$ the investor chooses the number of units χ_t that he buys of the risky asset so that he maximizes his utility from trading gains at the end of the year. More concretely, let the lottery L_t^T with probabilities $\{p_i\}$ and prizes $\{z_{i,t}^T\}$ represent the trading gains from time t to T . Then, at date t and state s_t (which determines the reference point), the investor with preference parameters θ solves

$$\max_{\chi_t} U(\theta, L_t^T, s_t)$$

subject to the constraints

$$\begin{aligned} W_{t'} &= (W_{t'-1} - \chi_{t'-1} P_{t'-1}) + \chi_{t'-1} P_{t'-1} R_{t'}, & \text{for } t' > t & \quad (\text{budget constraint}) \\ W_T &\geq 0 & & \quad (\text{no bankruptcy}) \\ \chi_t &\geq 0 & & \quad (\text{no shorting}) \end{aligned}$$

where U is the CPT utility functional defined in Section 4.1.³⁴ In particular, $U(\theta, L_t^T, s_t) = \sum_i w_i v(W_T - W_t^{\text{RP}})$ with $W_T = z_{i,t}^T + W_t$, and with the weights w_i , the value function v , and the reference wealth W^{RP} as defined in Equations 2 through 5. In particular, in this two-period model, the reference wealth level at $t = 0$ is equal to contemporaneous wealth, and at $t = 1$ it is equal to a convex combination between contemporaneous

³¹Using data on the risk-free (the one-month T-bill) rate and on stock returns from the Center for Research in Securities Prices for the period 1975–2011, we calculate that the average annual risk-free return and equity premium are 5.5% and 6.5%, and the average annual volatility of individual stocks is 64%, corresponding to semi-annual risk-free return $R_f = 1.027$ and annualized mean $\mu = 6.5\%$ and standard deviation $\sigma = 64\%$ for stock excess return. Calibrating the model to these values, our results below are even stronger.

³²Barber and Odean (2000) find that the mean/median U.S. household with a brokerage account holds 4.3/2.6 stocks and executes about 5 trades per year, corresponding to between 1 and 2 trades per stock per year. Alvarez, Guiso and Lippi (2012) find that the median number of yearly asset trades for Italian equity investors is 2.

³³Initial wealth W_0 can be arbitrarily chosen because having the same curvature for the value function over gains and losses implies that the value function is homogeneous, hence optimal actions do not depend on W_0 .

³⁴As in BX, the investor can borrow money. Since, in practice, most households do not borrow to buy stocks, solving the model *without* borrowing likely yields results that are more relevant to the empirical evidence, hence we also solve and present results for the case with a no-borrowing constraint, i.e., $\chi_t P_t - W_t \leq 0$.

wealth and the wealth at $t = 0$, i.e.,

$$\begin{aligned} W_0^{\text{RP}} &:= W_0 \\ W_1^{\text{RP}} &:= W_1 - \delta(W_1 - W_0), \end{aligned}$$

where δ is the memory decay parameter introduced in Equation 5, which determines how the reference point adjusts to incorporate gains and losses over time.

We solve the model numerically via backward induction, finding first the optimal actions $\chi_{1,d}^*$ and $\chi_{1,u}^*$ in the down ($R_1 = R_d$) and up ($R_1 = R_u$) state at $t = 1$, respectively, and then the optimal action χ_0^* at $t = 0$. Finding these optimal actions enables us to determine whether the investor participates in the stock market at $t = 0$ (i.e., $\chi_0^* > 0$) and whether he exhibits the DE (i.e., $\chi_{1,u}^* < \chi_0^* \leq \chi_{1,d}^*$). We solve the model for each of our individuals using the posterior mean estimates for his preference parameters, and we calculate the proportion that participates in the market and the proportion that exhibits the DE conditional on participation. Subsequently, we compare our results to those in BX, who solve this model for a representative agent who i) at $t = 1$ uses a reference point that perfectly sticks to the original wealth level at $t = 0$, which corresponds to the case $\delta = 1$; ii) has a value function with the Tversky and Kahneman (1992) median parameter estimates, i.e., $\alpha = 0.88$ and $\lambda = 2.25$; and iii) does not distort probabilities, which corresponds to the case $\gamma = \kappa = 1$.³⁵

Before we present our results, we note the following difference between the model of behavior in this application and in the estimation (see Section 4.1). In the former, the individual makes his choice at the beginning of the year taking into account his decision at midyear, while in the latter he makes his choice on each day in isolation. This approach corresponds to the “natural” or sensible frames in each setting: In the stock market, it is generally thought that individuals consider annual evaluation periods (see Benartzi and Thaler, 1995; Barberis, 2013), while in the betting market (and other similar markets that have been used for preference estimation; see Barseghyan et al., 2015) it seems most sensible that individuals narrowly frame each day’s choices between short-lived lotteries without looking ahead to the choices they might make the following week. Furthermore, this approach mirrors that of BX, who use the same model setup with the Tversky and Kahneman (1992) parameter estimates which are similarly derived from choice situations involving narrowly-framed short-lived lotteries.

³⁵BX set $\delta = 1$ because, in the absence of a previous estimate for δ , it is a sensible assumption that reflects the experimental evidence that prior outcomes affect the framing of subsequent prospects. For tractability reasons, they also assume no probability weighting.

Results In Table 7 we present, for various values of the equity risk premium μ , the proportions of individuals who optimally buy the stock at $t = 0$, and the proportions of individuals who exhibit the DE conditional on the initial purchase decision; we present results for the case with and without a no-borrowing constraint in Panels A and B, respectively. Like BX, we find that, as the risk premium increases, the proportion of individuals who buy the stock at $t = 0$ increases but the proportion of individuals who exhibit the DE decreases. The former result is straightforward while the explanation for the latter follows an intuition described above: The larger the potential gains relative to the potential losses, the likelier it is that after gains the individual moves to a region of the value function in which concavity is mild relative to loss aversion and so the prospect of holding on to the stock is more appealing than that of buying the stock in the first place. But while the trends that we find in behavior are as in BX, the magnitudes, hence our conclusions, are different. In particular, BX find that the (representative) agent does not buy the stock for any value of the equity premium μ up to 9%, and that he buys the stock but does not exhibit the DE for μ above 9%.³⁶ Instead, we find that a significant proportion of individuals participate in the stock market *and* exhibit the DE for a wide range of stock return parameter values. For example for $\mu = 6\%$, corresponding to a plausible equity premium and Sharpe ratio, we find that 65% (66%) of the individuals participate in the market at $t = 0$, and 57% (36%) of these individuals exhibit the DE in a model without (with) borrowing.³⁷ As a result, though BX raise valid concerns about the ability of prospect theory to explain the DE, our message is much more encouraging, since based on our estimates prospect theory *can* explain the DE for a much wider range of stock return parameter values.

Next, we examine why our results differ from those in BX. As mentioned above, the representative CPT agent in BX does not participate in the market for small values of the risk premium μ , and participates but does not exhibit the DE for large values of μ . Therefore, we focus on analyzing why individuals with our estimated parameters would participate for small μ and would exhibit the DE for both small and large μ .

³⁶BX propose a theory that is similar to prospect theory but in which utility is defined over *realized* gains and losses, and show that it can, to some extent, explain the DE (see also Barberis and Xiong, 2012 and Frydman et al., 2014). In this alternative model, the agent does not buy the stock for μ s up to 8%, exhibits the DE for $\mu \in [9\%, 11\%]$, and buys the stock but does not exhibit the DE for $\mu > 11\%$.

³⁷There is little empirical evidence on the prevalence of the DE in the population. Using a subsample of 7,965 individuals from the Barber and Odean (2000) data on the trading activity of 77,995 U.S. individuals, Dhar and Zhu (2006) calculate that 80% of the individuals exhibit the DE. Taking into account that this proportion is based on a measure — $DE := \text{Paper Gains Realized} - \text{Paper Losses Realized}$ — that overstates the prevalence of the DE relative to our measure, and that it is likely to be upward biased (as the authors note) due to the systematic exclusion of individuals from the original sample, the empirical evidence is close to our estimated proportion of around 60% of individuals exhibiting the DE for the case without borrowing and for μ corresponding to plausible combinations of the equity premium and Sharpe ratio.

Also, even though all preference parameters affect both the participation and the propensity to exhibit the DE, we focus on the most influential parameters for each behavior — loss aversion λ for participation and the memory decay parameter δ for the DE.

First, consistent with a substantial portion of the experimental literature, we estimate that a significant proportion (in fact, the majority) of individuals have a lower degree of loss aversion λ than the value of 2.25 used by BX.³⁸ These individuals are willing to buy the stock at $t = 0$ even for smaller values of the risk premium μ , for which they are more likely to exhibit the DE as we have argued previously. In Figure 9a, we plot the proportion of individuals who optimally participate in the stock market at $t = 0$ as a function of λ and μ , fixing the remaining parameters to our individual-specific estimates. We see, for example, that λ s below 2.0 are sufficient to get some participation for reasonable values of μ , and λ s below 1.7 are sufficient to get some participation for all μ s. Thus, allowing for heterogeneity in λ is mostly responsible for the higher participation rates, and therefore for our finding that prospect theory can explain the DE for low values of μ .

Second, we estimate that many individuals have a memory decay parameter δ significantly lower than the value of 1 used by BX. In Figure 9b, we plot the proportion of individuals who optimally exhibit the DE (conditional on participating) as a function of δ and μ , fixing the remaining parameters to our individual-specific estimates. We see that the relationship between δ and the proportion of individuals exhibiting the DE has an inverse-U shape. On the one end, $\delta = 0$ implies no DE since the optimal decision after gains and losses is the same; on the other end, $\delta = 1$ also implies no DE for large μ s, consistent with the insight of BX described earlier. However, intermediate δ s imply a higher prevalence of the DE for all μ s, because, after gains, the payoffs straddle the reference point in a way that loss aversion is compounded with risk aversion, hence the shape of the relevant part of the value function is even more concave (for an illustration of this intuition, see Figures 8a and 8c). For example, for $\mu = 10\%$ (the lowest risk premium for which there is participation, but no DE, with the parameters of BX) we find that a δ of 0.7 yields a DE rate of 66%, while a δ of 1 yields a DE rate of 6%. Thus, the intermediate values of the memory decay parameter δ that we have estimated for a significant fraction of individuals are mostly responsible for our finding that prospect theory can explain the DE for high values of μ .

³⁸We note that a low average value for λ is not necessary for obtaining a higher participation rate. Rather it is sufficient that there is substantial heterogeneity in λ across individuals such that a significant proportion, not necessarily a majority, has $\lambda < 2.25$ and therefore finds it optimal to buy the stock. Indeed, all experiments studied in the meta-analysis by Booij, van Praag and van de Kuilen (2010) find a large degree of variation in λ across individuals, some even finding significant proportions of individuals with λ smaller than 1.

6 Conclusion

In this paper, we develop a structural model of dynamic choice under risk within the prospect-theory paradigm, and estimate it using real-world data on individual trading activity in the sports wagering market. This market setting possesses attractive experimental-like features which facilitate the analysis of risk-taking behavior, but also attracts substantial amounts of money and shares many similarities with traditional financial markets. Our structural approach enables us to go beyond documenting an average relationship between prior outcomes and subsequent risk taking, to estimating the individual-level underlying economic primitives which give rise to this behavior. Our estimates extend support for existing experimental findings to a real market setting, but also reveal that individual risk attitude is characterized by substantial heterogeneity and history dependence. In particular, we find that the reference point individuals use to separate gains from losses depends on past profits and that this dependence varies considerably across individuals. Furthermore, we use our estimates in a model studying dynamic trading, to understand the implications of replacing the usual assumption of a representative prospect-theory agent who has the Tversky and Kahneman (1992) median preference parameters and a perfectly sticky reference point. Our estimates imply that prospect theory can explain the disposition effect for a much wider range of stock return parameter values than previously thought.

In light of our results, we conclude that it is important that future studies in finance incorporate heterogeneity and history dependence in individuals' risk preferences, and consider a wider range of preference parameter values than they currently do. For example, it would be interesting to further apply our findings to models that have used prospect theory to explain asset pricing irregularities, such as the equity premium and the volatility puzzles (e.g., Benartzi and Thaler, 1995; Barberis, Huang and Santos, 2001), and to examine whether the intuition of these studies continues to hold when we relax the assumption of a representative-agent with state-independent preferences.

References

- Abdellaoui, M, F Vossman, and M Weber.** 2005. "Choice-based elicitation and decomposition of decision weights for gains and losses under uncertainty." *Management Science*, 51(9): 1384–1399.
- Abdellaoui, M, H Bleichrodt, and C Paraschiv.** 2007. "Loss aversion under prospect theory: A parameter-free measurement." *Management Science*, 53(10): 1659–1674.
- Alvarez, F, L Guiso, and F Lippi.** 2012. "Durable consumption and asset management with transaction and observation costs." *American Economic Review*, 102(5): 2272–2300.
- Andrikogiannopoulou, A, and F Papakonstantinou.** 2016. "Estimating heterogeneous risk preferences from a panel of real-world betting choices." Unpublished Paper.
- Avery, C, and J Chevalier.** 1999. "Identifying investor sentiment from price paths: The case of football betting." *Journal of Business*, 72(4): 493–521.
- Bali, T. G, N Cakici, and R. F Whitelaw.** 2011. "Maxing out: Stocks as lotteries and the cross-section of expected returns." *Journal of Financial Economics*, 99(2): 427–446.
- Barber, B, and T Odean.** 2000. "Trading is hazardous to your wealth: The common stock investment performance of individual investors." *Journal of Finance*, 55(2): 773–806.
- Barber, B, and T Odean.** 2002. "Online investors: Do the slow die first?" *Review of Financial Studies*, 15(2): 455–487.
- Barberis, N.** 2013. "Thirty years of prospect theory in economics: A review and assessment." *Journal of Economic Perspectives*, 27(1): 173–96.
- Barberis, N, and M Huang.** 2008. "Stocks as lotteries: The implications of probability weighting for security prices." *American Economic Review*, 98(5): 2066–2100.
- Barberis, N, and W Xiong.** 2009. "What drives the disposition effect? An analysis of a long-standing preference-based explanation." *Journal of Finance*, 64(2): 751–784.
- Barberis, N, and W Xiong.** 2012. "Realization utility." *Journal of Financial Economics*, 104(2): 251–271.
- Barberis, N, M Huang, and T Santos.** 2001. "Prospect theory and asset prices." *Quarterly Journal of Economics*, 116(1): 1–53.
- Barseghyan, L, F Molinari, T O'Donoghue, and J. C Teitelbaum.** 2013. "The nature of risk preferences: Evidence from insurance choices." *American Economic Review*, 103(6): 2499–2529.
- Barseghyan, L, F Molinari, T O'Donoghue, and J. C Teitelbaum.** 2015. "Estimating risk preferences in the field." Unpublished Paper.
- Batagelj, V.** 1988. "Generalized Ward and related clustering problems." In *Classification and Related Methods of Data Analysis*, ed. H. H. Bock, 67–74. Amsterdam:North-Holland.
- Ben-Akiva, M, M Bergman, A. J Daly, and R Ramaswamy.** 1984. "Modeling inter-urban route choice behaviour." In *Proceedings of the 9th International Symposium on Transportation and Traffic Theory*, ed. Volmuller J. and Hamerslag R. Utrecht, The Netherlands:VNU Science Press.
- Benartzi, S, and R. H Thaler.** 1995. "Myopic loss aversion and the equity premium puzzle." *Quarterly Journal of Economics*, 110(1): 73–92.
- Benninga, S, and J Meishar.** 2000. "Heterogeneity and option pricing." *Review of Derivatives Research*, 4: 7–27.
- Bhamra, H. S, and R Uppal.** 2009. "The effect of introducing a non-redundant derivative on the volatility of stock-market returns when agents differ in risk aversion." *Review of Financial Studies*, 22(6): 2303–2330.
- Bollen, N. P, and R. E Whaley.** 2004. "Does net buying pressure affect the shape of implied volatility functions?" *Journal of Finance*, 59(2): 711–753.
- Booij, A, B van Praag, and G van de Kuilen.** 2010. "A parametric analysis of prospect theory's functionals for the general population." *Theory and Decision*, 68(1): 115–148.
- Boyer, B. H, and K Vorkink.** 2014. "Stock options as lotteries." *Journal of Finance*, 69(4): 1485–1527.

- Boyer, B, T Mitton, and K Vorkink.** 2010. "Expected idiosyncratic skewness." *Review of Financial Studies*, 23(1): 169–202.
- Brownstone, D, D. S Bunch, and K Train.** 2000. "Joint mixed logit models of stated and revealed preferences for alternative-fuel vehicles." *Transportation Research Part B: Methodological*, 34(5): 315–338.
- Cain, M, D Law, and D Peel.** 2000. "The favourite-longshot bias and market efficiency in UK football betting." *Scottish Journal of Political Economy*, 47(1): 25–36.
- Calvet, L. E, J. Y Campbell, and P Sodini.** 2009. "Fight Or flight? Portfolio rebalancing by individual investors." *Quarterly Journal of Economics*, 124(1): 301–348.
- Campbell, J. Y, and J. H Cochrane.** 1999. "By force of habit: A consumption-based explanation of aggregate stock market behavior." *Journal of Political Economy*, 107(2): 205–251.
- Campbell, J. Y, J. D Hilscher, and J Szilagyi.** 2011. "Predicting financial distress and the performance of distressed stocks." *Journal of Investment Management*, 9(2): 14–34.
- Chan, Y. L, and L Kogan.** 2002. "Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices." *Journal of Political Economy*, 110(6): 1255–1285.
- Chib, S, and E Greenberg.** 1995. "Understanding the Metropolis-Hastings algorithm." *The American Statistician*, 49(4): 327–335.
- Choi, S, R Fisman, D Gale, and S Kariv.** 2007. "Consistency and heterogeneity of individual behavior under uncertainty." *American Economic Review*, 97(5): 1921–1938.
- Cicchetti, C. J, and J. A Dubin.** 1994. "A microeconomic analysis of risk aversion and the decision to self-insure." *Journal of Political Economy*, 102(1): 169–186.
- Cohen, A, and L Einav.** 2007. "Estimating risk preferences from deductible choice." *American Economic Review*, 97(3): 745–788.
- Conrad, J, N Kapadia, and Y Xing.** 2014. "Death and jackpot: Why do individual investors hold overpriced stocks?" *Journal of Financial Economics*, 113(3): 455–475.
- Constantinides, G. M.** 1990. "Habit formation: A resolution of the equity premium puzzle." *Journal of Political Economy*, 98(3): 519–43.
- Coval, J. D, and T Shumway.** 2005. "Do behavioral biases affect prices?" *Journal of Finance*, 60(1): 1–34.
- Curcuro, S, D. J Lucas, J. C Heaton, and D Moore.** 2010. "Heterogeneity and portfolio choice: Theory and evidence." In *Handbook of Financial Econometrics: Tools and Techniques*, Vol. 1 of *Handbooks in Finance*, ed. Yacine Ait-Sahalia and Lars P. Hansen, 337–382. North-Holland.
- Dhar, R, and N Zhu.** 2006. "Up close and personal: Investor sophistication and the disposition effect." *Management Science*, 52(5): 726–740.
- Dorn, D, and P Sengmueller.** 2009. "Trading as entertainment?" *Management Science*, 55(4): 591–603.
- Durham, G. R, M. G Hertz, and J. S Martin.** 2005. "The market impact of trends and sequences in performance: New evidence." *Journal of Finance*, 60(5): 2551–2569.
- Fehr-Duda, H, and T Epper.** 2012. "Probability and risk: Foundations and economic implications of probability-dependent risk preferences." *Annual Review of Economics*, 4(1): 567–593.
- Feng, L, and M Seasholes.** 2005. "Do investor sophistication and trading experience eliminate behavioral biases in finance markets?" *Review of Finance*, 9(3): 305–351.
- Frazzini, A.** 2006. "The disposition effect and underreaction to news." *Journal of Finance*, 61(4): 2017–2046.
- Frydman, C, N Barberis, C Camerer, P Bossaerts, and A Rangel.** 2014. "Using neural data to test a theory of investor behavior: An application to realization utility." *Journal of Finance*, 69(2): 907–946.
- Gandar, J, R Zuber, T O'Brien, and B Russo.** 1988. "Testing rationality in the point spread betting market." *Journal of Finance*, 43(4): 995–1008.
- Gao Bakshi, X, and T.-C Lin.** 2015. "Do individual investors treat trading as a fun and exciting gambling activity? Evidence from repeated natural experiments." *Review of Financial Studies*, 28(7): 2128–2166.
- Garrett, T. A, and R. S Sobel.** 1999. "Gamblers favor skewness, not risk: Further evidence from United States lottery games."

- Economics Letters*, 63(1): 85–90.
- von Gaudecker, H, A van Soest, and E Wengstrom.** 2011. “Heterogeneity in risky choice behavior in a broad population.” *American Economic Review*, 101(2): 664–694.
- Gelman, A, J. B Carlin, H. S Stern, and D. B Rubin.** 1995. *Bayesian data analysis*. London, UK:Chapman and Hall.
- Geman, S, and D Geman.** 1984. “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images.” *Pattern Analysis and Machine Intelligence, IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6: 721–741.
- Gertner, R.** 1993. “Game shows and economic behavior: Risk-taking on ‘Card Sharks’.” *Quarterly Journal of Economics*, 108(2): 507–521.
- Glaser, M.** 2003. “Online broker investors: Demographic information, investment strategy, portfolio positions, and trading activity.” Unpublished Paper.
- Goeree, M. S.** 2008. “Limited information and advertising in the US personal computer industry.” *Econometrica*, 76(5): 1017–1074.
- Golec, J, and M Tamarkin.** 1998. “Bettors love skewness, not risk, at the horse track.” *Journal of Political Economy*, 106(1): 205–225.
- Gray, P. K, and S. F Gray.** 1997. “Testing market efficiency: Evidence from the NFL sports betting market.” *Journal of Finance*, 52(4): 1725–1737.
- Grinblatt, M, and M Keloharju.** 2001. “What makes investors trade?” *Journal of Finance*, 56(2): 589–616.
- Grinblatt, M, and M Keloharju.** 2009. “Sensation seeking, overconfidence, and trading activity.” *Journal of Finance*, 64(2): 549–578.
- Guiso, L, P Sapienza, and L Zingales.** 2013. “Time varying risk aversion.” Unpublished Paper.
- Hastings, W.** 1970. “Monte Carlo sampling methods using Markov chains and their applications.” *Biometrika*, 57(1): 97–109.
- Hens, T, and M Vlcek.** 2011. “Does prospect theory explain the disposition effect?” *Journal of Behavioral Finance*, 12(3): 141–157.
- Hey, J. D, and C Orme.** 1994. “Investigating generalizations of expected utility theory using experimental data.” *Econometrica*, 62(6): 1291–1326.
- Hoffman, A. O. I.** 2007. “Individual investors’ needs and the investment professional.” *Journal of Investment Consulting*, 8(2): 80–91.
- Jullien, B, and B Salanie.** 2000. “Estimating preferences under risk: The case of racetrack bettors.” *Journal of Political Economy*, 108(3): 503–530.
- Kliger, D, and O Levy.** 2009. “Theories of choice under risk: Insights from financial markets.” *Journal of Economic Behavior & Organization*, 71(2): 330–346.
- Kumar, A.** 2009. “Who gambles in the stock market?” *Journal of Finance*, 64(4): 1889–1933.
- Kuypers, T.** 2000. “Information and efficiency: An empirical study of a fixed odds betting market.” *Applied Economics*, 32(11): 1353–63.
- Kyle, A. S, H Ou-Yang, and W Xiong.** 2006. “Prospect theory and liquidation decisions.” *Journal of Economic Theory*, 129(1): 273–288.
- Lattimore, P, J Baker, and A Witte.** 1992. “The influence of probability on risky choice: A parametric examination.” *Journal of Economic Behavior & Organization*, 17(3): 377–400.
- Lemmon, M, and S. X Ni.** 2014. “Differences in trading and pricing between stock and index options.” *Management Science*, 60(8): 1985–2001.
- Levitt, S. D.** 2004. “Why are gambling markets organised so differently from financial markets?” *Economic Journal*, 114(495): 223–246.
- Liu, Y.-J, C.-L Tsai, M.-C Wang, and N Zhu.** 2010. “Prior consequences and subsequent risk taking: New field evidence from the Taiwan Futures Exchange.” *Management Science*, 56(4): 606–620.
- Li, Y, and L Yang.** 2013. “Prospect theory, the disposition effect, and asset prices.” *Journal of Financial Economics*, 107(3): 715–739.
- Manski, C. F.** 2004. “Measuring expectations.” *Econometrica*, 72(5): 1329–1376.
- Marschak, J.** 1960. “Binary-choice constraints and random utility indicators.” In *Stanford symposium on mathematical methods in*

- the social sciences*, Vol. 7, ed. K. J. Arrow, S. Karlin and P. Suppes, 19–38. Stanford, CA: Stanford University Press.
- Massa, M, and A Simonov.** 2005. “Behavioral biases and investment.” *Review of Finance*, 9(4): 483–507.
- McFadden, D.** 1974. “Conditional logit analysis of qualitative choice behavior.” In *Frontiers in Econometrics*, ed. P. Zarembka, 105–142. New York: Academic Press.
- Metrick, A.** 1995. “A natural experiment in ‘Jeopardy!’.” *American Economic Review*, 85(1): 240–253.
- Metropolis, N, A Rosenbluth, M Rosenbluth, A Teller, and E Teller.** 1953. “Equation of state calculations by fast computing machines.” *Journal of Chemical Physics*, 21(6): 1087–1092.
- Mitton, T, and K Vorkink.** 2007. “Equilibrium underdiversification and the preference for skewness.” *Review of Financial Studies*, 20(4): 1255–1288.
- Moskowitz, T. J.** 2015. “Asset pricing and sports betting.” Unpublished Paper.
- Moskowitz, T. J, and A Vissing-Jørgensen.** 2002. “The returns to entrepreneurial investment: A private equity premium puzzle?” *American Economic Review*, 92(4): 745–778.
- Ni, S. X.** 2008. “Stock option returns: A puzzle.” Unpublished Paper.
- Odean, T.** 1998. “Are investors reluctant to realize their losses?” *Journal of Finance*, 53(5): 1775–1798.
- Paravisini, D, V Rappoport, and E Ravina.** 2015. “Risk aversion and wealth: Evidence from person-to-person lending portfolios.” *Management Science*, forthcoming.
- Paul, R. J, and A. P Weinbach.** 2008. “Price setting in the NBA gambling market: tests of the Levitt model of sportsbook behavior.” *International Journal of Sport Finance*, 3: 137–145.
- Paul, R. J, and A. P Weinbach.** 2009. “Sportsbook behavior in the NCAA football betting market: Tests of the traditional and Levitt models of sportsbook behavior.” *Journal of Prediction Markets*, 3(2): 21–37.
- Polkovnichenko, V.** 2005. “Household portfolio diversification: A case for rank-dependent preferences.” *Review of Financial Studies*, 18(4): 1467–1502.
- Polkovnichenko, V, and F Zhao.** 2013. “Probability weighting functions implied in options prices.” *Journal of Financial Economics*, 107(3): 580–609.
- Pope, P. F, and D. A Peel.** 1989. “Information, prices and efficiency in a fixed-odds betting market.” *Economica*, 56(223): 323–41.
- Post, T, M. J van den Assem, G Baltussen, and R. H Thaler.** 2008. “Deal or no deal? Decision making under risk in a large-payoff game show.” *American Economic Review*, 98(1): 38–71.
- Rabin, M.** 2000. “Risk aversion and expected utility theory: A calibration theorem.” *Econometrica*, 68(5): 1281–1292.
- Ritter, J.** 1991. “The long-run performance of initial public offerings.” *Journal of Finance*, 46(1): 3–27.
- Shefrin, H, and M Statman.** 1985. “The disposition to sell winners too early and ride losers too long: Theory and Evidence.” *Journal of Finance*, 40(3): 777–790.
- Snowberg, E, and J Wolfers.** 2010. “Explaining the favorite-longshot bias: Is it risk-love or misperceptions?” *Journal of Political Economy*, 118(4): 723–746.
- Thaler, R. H, and E. J Johnson.** 1990. “Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice.” *Management Science*, 36(6): 643–660.
- Tversky, A, and D Kahneman.** 1992. “Advances in prospect theory: Cumulative representation of uncertainty.” *Journal of Risk and Uncertainty*, 5(4): 297–323.
- Wang, J.** 1996. “The term structure of interest rates in a pure exchange economy with heterogeneous investors.” *Journal of Financial Economics*, 41(1): 75–110.
- Weber, M, and C. F Camerer.** 1998. “The disposition effect in securities trading: An experimental analysis.” *Journal of Economic Behavior & Organization*, 33(2): 167–184.
- Woodland, L. M, and B. M Woodland.** 1994. “Market efficiency and the favorite-longshot bias: The baseball betting market.” *Journal of Finance*, 49(1): 269–79.
- Zhang, Y.** 2006. “Individual skewness and the cross-section of average stock returns.” Unpublished Paper.

Fixture and Event Description	Sport	Selection	Odds	Bet Type	Total Stake
Sachsenring 2010 (Moto GP, 8th round) <i>Winner</i>	Motorbikes	Rossi	3.20	Single	15
Chicago Fire vs Real Salt Lake (MLS) <i>Total Goals 2.5</i>	Soccer	Over 2.5	2.10	Singles	40
Lopez vs Concepcion (WBO) <i>Winner</i>	Boxing	Lopez	1.10		
Netherlands vs Spain (FIFA World Cup) <i>Match Winner</i>	Soccer	Spain	1.85	Treble	25
Seattle Mariners vs LA Angels (MLB) <i>Match Winner</i>	Baseball	Seattle	2.33		
Boston Red Sox vs Texas Rangers (MLB) <i>Total Match Runs Odd/Even</i>	Baseball	Odd	1.71		

39

Figure 1: Sample betting slip submitted by an individual. It contains three bet types: a “single” bet (a wager on one event), a “singles” bet (a wager on each of the possible “single” bets on the selected events), and a “treble” bet (a wager on three events).

Table 1: Summary Statistics of Individual and Bet Attributes

Summary statistics for the attributes of the individuals in our sample and their choices. Female is a dummy indicating gender. Age is in years. A bet day is a day during which an individual places a bet. Events per bet day is the mean number of events on which individuals bet in a bet day. Bet days per year per individual measures betting frequency.

	Mean	Median	Std. Dev.	Min	Max
<i>Individual Attributes:</i>					
Female	0.07	0	0.25	0	1
Age	32.85	31	9.63	18	67
<i>Bet Attributes:</i>					
Events per bet day	5.44	4	5.16	1	66
Bet days per individual	35.02	21	42.03	5	380
Bet days per year, per individual	64.76	41.87	63.93	1	327

Table 2: Summary Statistics of Lottery Characteristics


Summary statistics for the characteristics of the chosen day lotteries, pooled across all individuals. All monetary amounts are in euros.

	Percentiles						
	1st	5th	25th	50th	75th	95th	99th
Number of Prizes	2	2	2	2	6	61	354
Mean	-43.39	-17.54	-4.96	-2.05	-0.77	-0.13	-0.03
Standard Deviation	0.31	1.43	8.18	20.17	45.03	153.91	411.11
Skewness	-1.18	-0.32	0.72	2.49	5.94	33.23	180.94
Kurtosis	1.00	1.26	2.77	8.16	28.56	265.59	$4 \cdot 10^4$
Bet Amount	0.10	0.55	3.90	10.00	25.00	100.00	311.68
Maximum Prize	0.33	2.02	21.19	79.08	276.08	$3 \cdot 10^3$	$2 \cdot 10^4$
Minimum Probability	$1 \cdot 10^{-10}$	$3 \cdot 10^{-6}$	$3 \cdot 10^{-3}$	0.03	0.15	0.44	0.49
Maximum Probability	0.06	0.17	0.54	0.78	0.94	1.00	1.00


Table 3: Reduced-form Analysis

Panel A presents results from a panel OLS regression in which the dependent variable is the logarithm of the duration (in days) of each “no play” event — i.e., the length of the interval between consecutive play days — for each individual. Panel B presents results from panel OLS regressions in which the dependent variables are the logarithms of appropriate affine transformations of the mean, variance, and skewness (specifically, we multiply the mean by -1 and add a constant to the skewness to ensure they are positive) of the payoffs of the lottery chosen by each individual on each play day, and the logarithm of the monetary stake wagered on that lottery. The explanatory variables are linear and quadratic terms of past profits, measured as cumulative gains/losses (in hundreds of euros) realized by each individual over the 7 calendar days preceding the beginning of the “no play” event in Panel A (the play day in Panel B). Individuals with 5 or fewer observations are excluded from the analysis. All specifications include individual-specific fixed effects. Robust t -statistics are reported below the coefficients. $^*/^{**}/^{***}$ indicate significance at the 10% /15% /1% levels.


	Panel A: Participation	Panel B: Lottery Choice			
	Duration	Mean	Variance	Skewness	Stake
Constant	0.70 ***	0.37 ***	2.28 ***	0.81 ***	2.47 ***
	4.61	2.57	22.79	24.85	23.81
CumGain	-0.08 ***	0.14 ***	0.12 ***	-0.01	0.14 ***
	-6.23	9.05	9.12	-0.87	8.69
CumLoss	-0.03 ***	0.09 ***	0.09 ***	0.02 **	0.08 ***
	-3.02	6.74	7.30	2.07	6.31
CumGain ²	0.01 ***	-0.00 ***	-0.00 ***	0.00	-0.00 ***
	4.41	-7.27	-7.67	1.00	-7.23
CumLoss ²	0.00 **	-0.00 ***	-0.00 ***	-0.00 **	-0.00 ***
	2.38	-6.26	-6.92	-2.18	-6.11
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Obs	11,020	11,332	11,332	11,332	11,332
Adjusted R^2	15.85%	48.68%	52.94%	45.47%	51.38%




(a) Probability weighting for different values of γ .



(b) Density weighting for different values of γ .



(c) Probability weighting for different values of κ .



(d) Density weighting for different values of κ .

Figure 2: The Lattimore, Baker and Witte (1992) probability weighting function (in Panels *a* and *c*) and its derivative (in Panels *b* and *d*) for several values of the curvature parameter γ and the elevation parameter κ . In Panels (a) and (b), we present plots for $\kappa = 1$ and several values of $\gamma \leq 1$, and in Panels (c) and (d) we present plots for $\gamma = 1$ and several values of $\kappa \geq 1$. In CPT, a lottery's utility equals the expectation of the value (according to the value function) of the outcome (relative to the reference point) with respect to a transformation (according to the probability weighting function) of the outcome's distribution. The probability weighting function w transforms the cumulative (over losses) or the decumulative (over gains) distribution function. Its derivative w' is the weighting factor applied to the density — for continuous distributions — or, roughly, the probability per unit interval — for discrete distributions; the plots of w' in Panels (b) and (d) show, for positive prizes, how this weighting factor depends on the location of the outcome in the distribution.

Table 4: Posterior Estimates of Population Mean and Variance

This table presents summary statistics of the posterior estimates for the population mean (in Panel A) and variance (in Panel B) of the elements of ϑ_n . The 95% HPDI is the smallest interval such that the posterior probability that a parameter lies in it is 0.95. NSE stands for autocorrelation-adjusted numerical standard errors for the posterior mean estimate of each parameter.

	Panel A: Means					Panel B: Variances				
	Mean	Median	Std.Dev.	95% HPDI	NSE	Mean	Median	Std.Dev.	95% HPDI	NSE
π	0.24	0.24	0.01	[0.22 , 0.27]	0.00	0.05	0.05	0.00	[0.04 , 0.06]	0.00
τ	1.36	1.35	0.11	[1.15 , 1.58]	0.01	2.70	2.65	0.39	[2.06 , 3.53]	0.04
α	0.86	0.86	0.01	[0.85 , 0.87]	0.00	0.01	0.01	0.00	[0.01 , 0.01]	0.00
λ	1.48	1.48	0.04	[1.39 , 1.56]	0.01	0.25	0.24	0.03	[0.20 , 0.32]	0.01
γ	0.88	0.88	0.01	[0.87 , 0.90]	0.00	0.01	0.01	0.00	[0.01 , 0.01]	0.00
κ	1.41	1.41	0.06	[1.31 , 1.53]	0.01	0.68	0.61	0.18	[0.46 , 1.13]	0.04
δ	0.60	0.60	0.02	[0.55 , 0.64]	0.00	0.15	0.15	0.00	[0.14 , 0.16]	0.00

Table 5: Estimated Percentiles of Model Parameters

This table presents the 5th, 25th, 50th, 75th, and 95th percentiles of the estimated population distributions, for all model parameters.

	Percentiles				
	5 th	25 th	50 th	75 th	95 th
π	0.02	0.07	0.15	0.34	0.71
τ	0.08	0.31	0.69	1.50	4.58
α	0.73	0.83	0.87	0.90	0.95
λ	0.92	1.18	1.37	1.70	2.35
γ	0.75	0.85	0.91	0.94	0.98
κ	0.64	1.00	1.17	1.59	2.85
δ	0.02	0.34	0.66	0.87	0.99




Figure 3: Estimated population densities of elements of ϑ_n : the probability π of having the opportunity to place a bet on any given day, the inverse scale τ of the random choice errors, the measures α and λ of curvature and loss aversion of the value function, the measures γ and κ of curvature and elevation of the probability weighting function, and the memory decay parameter δ . The blue solid line plots our baseline estimated densities, and the black dotted and red dash-dotted lines plot them for a low-variance prior ($\underline{\Lambda} = 0.1I$) and a high-variance prior ($\underline{\Lambda} = 10I$), where $\underline{\Lambda}$ is the hyper-parameter of the prior distribution of $\Sigma_{\tilde{\vartheta}}$ (see Equations 7 and 10 in Section 4.2).




Figure 4: Value function and probability weighting function for the CPT specification, for various combinations of selected percentiles from the estimated distribution of parameters α , λ , γ , and κ . In the left (right) panel, the black solid lines correspond to the medians of parameters α and λ (γ and κ), and the remaining lines correspond to combinations of a low value — the 25th percentile — and a high value — the 75th percentile — of each parameter with the median value of the other parameter.




Figure 5: Illustration of goodness of fit of our model. We plot the histogram of the mean — across all choices for each individual — ratio ranking of the chosen lotteries according to the deterministic utility component. For choices that are more preferred, the ranking is close to 0, and for choices that are less preferred it is close to 1. For example, the 5th most-preferred out of 100 alternatives has ratio ranking equal to 0.05.

Table 6: Comparison of Our Model with Alternatives

Mean and median log likelihood, across individuals, and Deviance Information Criterion (DIC) for our model and four alternatives: i) one in which lotteries are chosen randomly; ii) one in which individuals are assumed to be risk neutral, iii) one in which all individuals' preferences are neither heterogeneous nor history-dependent, and iv) one in which individuals' preferences are heterogeneous but not history-dependent. The DIC is calculated as $\overline{D(\vartheta)} + p_D$, where $D(\vartheta) = -2 \log p(x, y | \vartheta)$, $\overline{D(\vartheta)} = E_{\vartheta} [D(\vartheta) | x, y]$, and $p_D = \overline{D(\vartheta)} - D(\bar{\vartheta})$, with $\bar{\vartheta} = E[\vartheta | x, y]$, i.e., the posterior mean.

Model	log L		DIC
	Mean	Median	
Random Choice	-251.76	-150.51	503.52
Risk Neutrality	-217.55	-130.08	435.10
No Heterogeneity, No History dependence	-210.94	-121.19	421.88
Heterogeneity, No History dependence	-179.27	-106.98	362.10
Our Model	-178.13	-105.26	356.80




Figure 6: Plots of the mean excess realized returns for each percentile group of odds in our fixed-odds sports wagering market and in the parimutuel horse-racetrack betting market. For our market, we plot the mean excess realized return for each odds group in a blue solid line and the 95% bootstrap confidence intervals in blue dotted lines. For the parimutuel market, we reconstruct the plot in Figure 1 of Snowberg and Wolfers (2010) in a red dashed line. The mean excess realized return is calculated in excess of the mean realized return across all odds groups in the market.




Figure 7: Illustration of the robustness of the estimation results to alternative assumptions for the subjective beliefs. We plot the estimated population densities of elements of ϑ_n : the probability π of having the opportunity to place a bet on any given day, the inverse scale τ of the random choice errors, the measures α and λ of curvature and loss aversion of the value function, the measures γ and κ of curvature and elevation of the probability weighting function, and the memory decay parameter δ . In both panels, the blue solid line corresponds to our baseline estimated densities. In Panel (a), the red dash-dotted lines correspond to the estimated densities under the assumption that the probabilities individuals associate with bet outcomes equal the win frequencies of past outcomes with similar prices. In Panel (b), the black dotted and red dash-dotted lines correspond to estimations that assume subjective beliefs exhibit small and large deviations, respectively, from the rational beliefs as approximated by the odds-implied probabilities in our baseline analysis. For more details on the alternative specifications, see Section 4.5.





Figure 8: Illustration of the behavior of three benchmark individuals toward a lottery in two hypothetical states: one in which they have no previous profits (neutral state), and one in which they have previously realized a gain (gain state). The estimated preference parameters for Individual A are $\alpha = 0.88$, $\lambda = 1.41$, $\delta = 1$, $\gamma = \kappa = 1$; for Individual B they are $\alpha = 0.89$, $\lambda = 1.43$, $\delta = 0.5$, $\gamma = \kappa = 1$, and for Individual C they are $\alpha = 0.50$, $\lambda = 1.12$, $\delta = 1$, $\gamma = \kappa = 1$. The lottery under consideration pays -7 or 10 with equal probability. In Panel (a) Individual A faces the lottery without prior gains and after a large prior gain of 8 , in Panel (b) Individual A faces the lottery without prior gains and after a small prior gain of 4 , in Panel (c) Individual B faces the lottery without prior gains and after a large prior gain of 8 , and in Panel (d) Individual C faces the lottery without prior gains and after a large prior gain of 8 . In all panels, the blue curves plot the individual's value function; the black dots denoted by L_1 and G_1 indicate the values corresponding to the negative and positive payoff, respectively, from the lottery at zero level of past profits; the magenta x-mark denoted by CE_1 indicates the certainty equivalent of the lottery; and the black dot denoted by O_1 indicates the value (and certainty equivalent) of rejecting the lottery. Also, the green dots denoted by L_2 and G_2 indicate the values (as perceived by each individual) corresponding to the negative and positive payoff, respectively, from the lottery after prior gains; the magenta x-mark denoted by CE_2 indicates the certainty equivalent of the lottery; and the green dot denoted by O_2 indicates the value (and certainty equivalent) of rejecting the lottery.


Table 7: Disposition Effect

This table presents results on stock-market participation and the disposition effect for the two-period portfolio problem presented in Section 5. Panel A presents results for the case in which borrowing is *not* allowed, and Panel B presents results for the case in which borrowing *is* allowed. Solving the model for each individual in the sample using our posterior mean estimates of his preference parameters (see Section 4), in each panel we present, for several values of the equity risk premium μ (expressed as an annual percentage), the proportions of individuals who are willing to buy the stock at $t = 0$ and the proportions of individuals who exhibit the disposition effect (DE) conditional on having bought the stock at $t = 0$. An individual exhibits the DE if he optimally decreases (increases or maintains) his stock position at $t = 1$ after realizing gains (losses) in the first period.

μ	Panel A : Without Borrowing		Panel B : With Borrowing	
	Stock Buyers	Conditional DE	Stock Buyers	Conditional DE
3%	36.61%	65.04%	35.42%	44.54%
4%	49.40%	71.08%	50.00%	44.64%
5%	58.04%	62.05%	57.74%	38.14%
6%	65.48%	57.27%	66.37%	36.32%
7%	74.11%	46.99%	76.19%	32.42%
8%	82.14%	42.39%	83.93%	28.37%
9%	87.20%	41.30%	88.69%	22.48%
10%	89.58%	35.55%	89.58%	12.96%
11%	91.96%	27.83%	91.96%	10.36%
12%	91.96%	22.98%	92.26%	5.48%
13%	92.56%	14.47%	92.86%	4.49%



(a) Proportion participating.



(b) Proportion exhibiting disposition effect.

Figure 9: Results on stock-market participation and the disposition effect for the two-period portfolio problem of Section 5, for the case without borrowing. Panel (a) plots the proportion of individuals who participate (i.e., $\chi_0^* > 0$) as a function of loss aversion λ and the equity risk premium μ (expressed as an annualized percentage). Panel (b) plots the proportion of individuals exhibiting the disposition effect (i.e., $\chi_{1,u}^* < \chi_0^* \leq \chi_{1,d}^*$) conditional on participating, as a function of the memory decay parameter δ and the equity risk premium μ .

Appendix

A Gibbs sampler

To estimate the model in Section 4, we use Bayesian methods, which involve deriving the joint posterior distribution of all model parameters conditional on the observed choices. The joint posterior of $\mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}}, \{\vartheta_n\}$ is proportional to the product of the likelihood in Equation 8 and the joint prior in Equation 9, but due to the model's complexity we cannot calculate it analytically, therefore we utilize the Markov Chain Monte Carlo (MCMC) algorithm to make draws from it. In particular, we use the Gibbs sampler (Geman and Geman, 1984), according to which we partition the model parameters into three blocks ($\mu_{\tilde{\vartheta}}, \Sigma_{\tilde{\vartheta}},$ and $\{\vartheta_n\}$), and in each iteration of the algorithm we sequentially draw from the conditional posterior of one block given the data and the draws for the other blocks from the previous iteration. The resulting sequence of draws is a Markov chain with a stationary distribution that converges to the joint posterior.

In the first iteration, $k=0$, we pick starting values for the model parameters by randomly drawing from their priors; estimation results should not depend on the starting values if the algorithm explores the posterior, and indeed they are essentially unchanged if we start the estimation at different values. Given values of the parameters for each individual and for their population mean and variance in iteration k , in iteration $k+1$ we perform the following steps:

1. Given values $\{\vartheta_n^{(k)}\}$ of the model parameters for each individual and value $\Sigma_{\tilde{\vartheta}}^{(k)}$ for their population variance, we draw a value $\mu_{\tilde{\vartheta}}^{(k+1)}$ for their population mean. In particular, using a normal prior with mean $\underline{\kappa}$ and variance \underline{K} for $\mu_{\tilde{\vartheta}}$, its conditional posterior is $\mathcal{N}(\bar{\kappa}, \bar{K})$, i.e., a normal with mean $\bar{\kappa} := \bar{K} \left(\underline{K}^{-1} \underline{\kappa} + \Sigma_{\tilde{\vartheta}}^{-1} \sum_{n=1}^N \tilde{\vartheta}_n \right)$ and variance $\bar{K} := \left(\underline{K}^{-1} + \Sigma_{\tilde{\vartheta}}^{-1} N \right)^{-1}$, where N is the number of individuals, and $\tilde{\vartheta}_n := g(\vartheta_n)$ are transformations of ϑ_n (presented in Section 4) such that $\tilde{\vartheta}_n \in (-\infty, +\infty)$. Setting $\underline{K} = 100I$ makes the prior weak and renders $\underline{\kappa}$ immaterial (so we set $\underline{\kappa} = \mathbf{0}$).
2. Given values $\{\vartheta_n^{(k)}\}$ of the model parameters for each individual and value $\mu_{\tilde{\vartheta}}^{(k+1)}$ for their population variance, we draw a value $\Sigma_{\tilde{\vartheta}}^{(k+1)}$ for their population variance. Using a Wishart prior with $\underline{\lambda}$ degrees of freedom and scale matrix $\underline{\Lambda}^{-1}$ for $\Sigma_{\tilde{\vartheta}}^{-1}$, its conditional posterior is $\mathcal{W}(\bar{\lambda}, \bar{\Lambda}^{-1})$, i.e., a Wishart with $\bar{\lambda} := \underline{\lambda} + N$ degrees of freedom and scale matrix $\bar{\Lambda}^{-1} := \left(\underline{\Lambda} + \sum_{n=1}^N (\tilde{\vartheta}_n - \mu_{\tilde{\vartheta}}) (\tilde{\vartheta}_n - \mu_{\tilde{\vartheta}})' \right)^{-1}$.

Setting $\underline{\lambda} = \text{rank}(\Sigma_{\tilde{\vartheta}})$ (which is the minimum for a proper prior) makes the prior weak, while for $\underline{\Lambda}$ we find that estimation results are not sensitive to using $\underline{\Lambda} = 0.1I$, $\underline{\Lambda} = I$, or $\underline{\Lambda} = 10I$.

3. Given the data for each individual and values $\mu_{\tilde{\vartheta}}^{(k+1)}$ and $\Sigma_{\tilde{\vartheta}}^{(k+1)}$ for the population mean and variance of the model parameters, we draw values $\{\vartheta_n^{(k+1)}\}$ for the parameters for each individual. Since we do not have a conjugate prior for the individual-specific parameters, the conditional posterior is not from a known distribution family, rather we only know its kernel, i.e., that it is proportional to the prior times the likelihood (see Equations 7 and 8), and therefore we need to draw from it using the Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970). The idea behind the algorithm is to make draws from a candidate generating density, and then accept them or not, balancing their posterior density against their generating density, in an effort to explore the parameter space but also to concentrate on areas of high posterior probability. We generate the candidate draw ϑ'_n from $\vartheta_n^{(k)}$ using a Gaussian random walk on the transformed parameters $\tilde{\vartheta}_n^{(k)}$; that is, we use $\vartheta'_n := g^{-1}(g(\vartheta_n^{(k)}) + \sigma L\eta)$, where $\eta \sim \mathcal{N}(0, I)$, $\Sigma_{\tilde{\vartheta}}^{(k+1)} = LL'$ is the Cholesky decomposition, and $\sigma > 0$ is chosen so that the “jumps” $\sigma L\eta$ are neither too small (hence we accept too many draws) nor too large (hence we accept too few draws), to enable us to properly explore the whole posterior (Gelman et al., 1995). Letting $q(\cdot)$ be the density of this generating process, we subsequently use the acceptance probability $\alpha(\vartheta_n; \vartheta'_n) := \min\{p(\vartheta'_n|\cdot)q(\vartheta_n) / p(\vartheta_n|\cdot)q(\vartheta'_n), 1\}$ calculated by Chib and Greenberg (1995) to either accept the draw and set $\vartheta_n^{(k+1)} = \vartheta'_n$, or to reject it and set $\vartheta_n^{(k+1)} = \vartheta_n^{(k)}$.

B Generation of alternatives

To randomly generate a lottery from the sportsbook, we use the procedure that people essentially follow when they select a day lottery, i.e., we draw i) the number of bets in the lottery, ii) the type, money risked, and number of individual wagers for each bet, and iii) the winning odds for each individual wager. All these characteristics are drawn from parametric approximations (estimated by maximum likelihood) of their empirical distributions in our data. While the winning odds of the individual wagers are not correlated with the other lottery characteristics in our data, the number of bets in the day lottery and the bet type, the money risked, and the number of individual wagers show strong dependence, as expected, hence we draw them jointly.

First, we draw the number of bets (*numBets*) in a lottery from a negative binomial distribution fitted to the number of bets per day lottery across all day lotteries chosen by individuals in our data. Second, we draw

the bet type. Bet types can be broadly classified into two categories — *permutation* and *full-cover* bets — so we first draw the category using a Bernoulli distribution fitted to their frequency. Permutation bets include all accumulators of (some) type k that can be constructed from a given number of wagers, hence can be indexed by k , so conditional on drawing the permutation category, we draw k from a negative binomial fitted to the indices of permutation bets in the observed day lotteries with $numBets$ number of bets. Full-cover bets include all accumulators (or all except for singles) that can be constructed from a given number of wagers; such bets are few and rarely selected, so conditional on drawing the full-cover category, we draw the specific full-cover bet type from a multinomial with event probabilities equal to their relative proportions in all day lotteries in the data. Third, we draw the number of wagers ($numWagers$) involved in a bet. In some cases, the bet type drawn determines uniquely the number of wagers, while in others it determines the minimum number of wagers. As a result, given the drawn bet type, we draw the number of wagers from a negative binomial truncated to the appropriate support and fitted to the number of wagers involved in all bets in our data belonging to the drawn bet type. Fourth, we draw the amount of money risked, which is divided evenly among the number of accumulators in the bet. Conditional on drawing a permutation bet, we draw the amount from a log-normal with mean linear in the index k that is fitted to the amounts staked in all permutation bets in the data. Conditional on drawing a full-cover bet, we draw the amount from a log-normal fitted to the amounts staked in all bets in our data of the specific full-cover bet type drawn. Finally, we draw the odds for each of the $numWagers$ wagers, by drawing the commission from a Johnson distribution fitted to the commissions in all individual wagers we observe, and then drawing the implied win probability from a uniform between the minimum and the maximum implied probability we observe.

Using this procedure, we randomly draw 20,000 lotteries, which we then group into 100 clusters.³⁹ We do so using the generalized Ward hierarchical agglomeration algorithm (Batagelj, 1988), which starts by placing each lottery in a separate cluster, and in each subsequent step combines two clusters such that it minimizes the mean distance between lotteries and the center of their cluster, where we use the Wasserstein distance (defined, for densities f_X and f_Y , as $\inf_{f_{XY}} E[\|X - Y\|]$, where the infimum is over all joints with marginals f_X, f_Y). After grouping lotteries into clusters, from each cluster we choose the lottery that is most representative, i.e., the one with the minimum mean distance to other lotteries in the cluster. The resulting set of lotteries should reasonably “span” the set of randomly drawn lotteries.

³⁹The amount of computer memory required to group lotteries into clusters increases quadratically with the number of lotteries, hence the number of lotteries we draw randomly is constrained by memory size.