King's Research Portal

DOI:

[10.1093/rfs/hhz127](https://doi.org/10.1093/rfs/hhz127)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Andrikogiannopoulou, A., & Papakonstantinou, F. (2020). History-Dependent Risk Preferences: Evidence from Individual Choices and Implications for the Disposition Effect. *REVIEW OF FINANCIAL STUDIES*, 33(8), 3674–3718. <https://doi.org/10.1093/rfs/hhz127>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

History-Dependent Risk Preferences: Evidence from Individual Choices and Implications for the Disposition Effect

Angie Andrikogiannopoulou
King's College London

Filippos Papakonstantinou
King's College London

Using trading data from a sports wagering market, we estimate individuals' dynamic risk preferences within a prospect theory paradigm. This market's experimental-like features facilitate preference estimation, and our long panel enables us to study whether preferences vary across individuals and depend on earlier outcomes. Our estimates extend support for experimental findings—mild utility curvature, moderate loss aversion, and probability overweighting of extreme outcomes—to a market setting and reveal that preferences are heterogeneous and history dependent. Applying our estimates to a portfolio choice problem, we show prospect theory can better explain the prevalence of the disposition effect than previously thought. (*JEL* D03, D81, G02, G11)

Received XXXX XX, XXXX; editorial decision XXXX XX, XXXX by Editor
XXXXXXXXXXXXXX.

Though a large body of experimental evidence supports prospect theory, its more widespread application and acceptance in finance is hindered

We are grateful to Francesca Cornelli (the editor) and two anonymous referees for very helpful comments. We also thank Nick Barberis, Kenza Benhima, Adrian Bruhin, Thomas Epper, Harald Hau, Bo Honore, Alex Kostakis, Andriy Norets, Kim Peijnenburg, Chris Rogers, Bernard Salanie, Luis Santos-Pinto, Hyun Shin, Dimitri Vayanos, and Wei Xiong; participants at the 2015 Miami Behavioral Finance Conference and the 2016 CEPR European Conference on Household Finance; and seminar participants at Athens University of Economics and Business, Cambridge University, Cass Business School, Imperial College London Business School, Goethe University Frankfurt, 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 for their advice. This work was supported by an Amazon Web Services in Education research grant, and by a Swiss Finance Institute grant. Send correspondence to Filippos Papakonstantinou, King's Business School, Bush House, 30 Aldwych, London WC2B 4BG, UK; telephone: +44 (0) 20 7848 3770. E-mail: fpapakon@kcl.ac.uk.

doi:10.1093/rfs/Sample

Advance Access publication XXXXXXXX XX, XXXX

by the lack of nonexperimental evidence on choice under risk. On the one hand, preference estimates from experiments face objections regarding their generalizability to market settings involving real money. On the other hand, traditional financial markets do not allow for a clean identification of preferences, due to the complex structure of asset returns. Most assets do not have an exogenous settling-up point, and returns are correlated with aggregate risks; therefore, return distributions cannot be accurately predicted or estimated.

In this paper, we use the sports wagering market as a real-world financial market laboratory to estimate individual preferences within the prospect theory paradigm. Assets in this market (i.e., wagers) bear no systematic risk, are short lived, and have a terminal payoff exogenously determined by match outcomes. Furthermore, a wager's quoted price is often a good predictor of its win probability. Thus, assets' payoff distributions are observable with significant accuracy. This allows for a lottery representation of individual choices and facilitates preference identification. At the same time, sports wagering markets share similarities to financial markets, both in their organization and in their participants' mentality. Thus, they constitute an ideal setting for preference estimation that can yield valuable insights about financial markets.¹

In a preliminary analysis of behavior in this market, we find that past profits affect choices in a manner consistent with the house money and break-even effects (Thaler and Johnson 1990): after gains, individuals increase their betting frequencies, wager higher amounts, and select lotteries with higher variance, while after losses they favor lotteries with higher skewness. This behavior is at odds with the traditional expected utility theory in which past outcomes affect future choices only through a wealth effect, and naturally lends itself to a reference-dependent theory, in which past outcomes may affect the reference point and hence subsequent behavior.

Motivated by this evidence, we model risk-taking using an extension of the Tversky and Kahneman (1992) cumulative prospect theory (CPT), the most prominent reference-dependent theory of choice. In CPT, individuals (1) evaluate risks by computing gains and losses relative to a reference point; (2) use a value function that is concave for gains and convex for losses, corresponding to risk aversion and risk seekingness, respectively; (3) have higher sensitivity to losses than gains, corresponding to loss aversion; and (4) systematically distort probabilities. As is common in the literature, we also assume individuals

¹ The literature has long identified the relevance of the sports wagering markets for traditional financial markets (e.g., Pankoff 1968; Durham, Hertz, and Martin 2005; Moskowitz 2018), but has been confined to aggregate analyses based on market price data.

frame choices narrowly, evaluating risks in one context separately from other risks. To incorporate history dependence and extend CPT to dynamic situations, we propose that the reference point may adapt to wealth fluctuations sluggishly. Inspired by the standard exponential-discounting assumption, we model the reference point as a function that attaches exponentially lower weights to earlier outcomes. Essentially, the discount factor captures the rate of memory decay which measures how far back each individual recalls past outcomes. Thus, the reference point may update immediately, in which case it equals contemporaneous wealth, or it may update slowly, in which case it “sticks” between past and current wealth, giving rise to time-varying risk preferences. This parsimonious approach has the advantage of introducing preference history dependence through a single intuitive parameter, which can be estimated and used in applications. To estimate our model of behavior, we embed it in an econometric model in which individual preferences are heterogeneous and drawn from a population distribution.

We find that, on average, individuals use a value function that is concave (convex) over gains (losses), exhibit loss aversion, and overweight extreme positive and negative outcomes. Thus, our results validate the general features of individual risk-taking behavior estimated in experiments, and show that they carry over to real-world market settings. This is significant, given that an increasing number of studies appeal to prospect theory to explain various behaviors and phenomena observed in financial markets (see the Barberis 2018 review). Though our parameter estimates are well within the range of experimental estimates, our estimate of loss aversion lies toward the lower end of the range and is substantially lower than the widely used estimate from the Tversky and Kahneman (1992) experiment. This may be because laboratory results are not transferable to naturally occurring market environments like the one we observe (see, e.g., Harrison, List, and Towe 2007), or possibly because more loss-averse people might not participate in a market with negative expected returns. The latter would indicate that our results are less relevant for the general population, but possibly more relevant for finance, because financial market participants may be selected in a similar way; that is, more loss-averse people may stay out of both markets.² Other similarities between wagering markets and traditional financial markets, which we discuss below, further suggest that our results may be particularly relevant for finance. As such, our low estimate of loss aversion is intriguing given that a high loss aversion has found mixed success in explaining puzzles in finance. For example,

² Models (e.g., Barberis, Huang, and Thaler 2006) show that loss aversion explains the observed low stock market participation rate, and evidence suggests that more loss-averse people are less likely to participate in the stock market (Dimmock and Kouwenberg 2010).

although high loss aversion can help explain the low stock market participation rate and the equity premium puzzle, it counterfactually implies higher mean returns for high-beta stocks and for stocks with high idiosyncratic volatility and low prevalence of the disposition effect, that is, individuals' tendency to sell (retain) stocks whose value has increased (decreased) since purchase.

We also find that risk preferences are widely heterogeneous across individuals and history dependent, as the reference point most individuals use to separate gains from losses sticks to previous levels. The distribution of the memory decay parameter reveals that past outcomes affect individuals to varying degrees, with a marked concentration close to the two natural extremes of no memory versus perfect memory. These findings can provide justification as well as guidance to the growing literature that appeals to preference heterogeneity and time variation to explain a variety of behaviors and anomalies in finance. For example, preference heterogeneity is useful for explaining the observed heterogeneity in households' portfolio choices, from, for example, stock market participation to portfolio diversification (for a review, see Curcuru et al. 2010), and also observed patterns in trading volume, the equity premium puzzle, and equilibrium asset prices.³ Furthermore, time-varying preferences are useful for understanding why households' portfolio choices vary with wealth fluctuations and macroeconomic shocks and for explaining trading patterns, such as the disposition effect. 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 countercyclicality of the equity premium and the excess volatility of stock returns.⁴

In contrast to our estimation results, most applied behavioral models in finance assume (a) a representative agent, (b) the median CPT parameter estimates of Tversky and Kahneman (1992), and (c) that individuals have either no memory or perfect memory. In an application, we consider some implications of our preference estimates on dynamic trading, specifically on the disposition effect. Barberis and Xiong (2009) consider a prospect theory agent with perfect memory and show that, contrary to conventional wisdom, prospect theory does not provide a plausible explanation of the disposition effect: the agent does not invest in stocks if the risk premium is low, whereas if the risk premium is high, he is more inclined to sell stocks with prior losses than with prior gains; that is, he exhibits the opposite of the disposition effect.

³ See, for example, Wang (1996), Benninga and Meishar (2000), Chan and Kogan (2002), and Bhamra and Uppal (2009).

⁴ See, for example, Constantinides (1990), Campbell and Cochrane (1999), and Barberis, Huang, and Santos (2001).

Though this is an intriguing insight, the extent to which it “bites” in reality depends on people’s actual preference parameters. Applying our estimated preferences to a portfolio problem similar to Barberis and Xiong (2009), we encouragingly find that prospect theory *can* explain the prevalence of the disposition effect, for a wide range of risk premiums. Specifically, for the observed 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, which is roughly in line with empirical observations. The difference in our findings can be attributed to the fact that we estimate (a) substantial preference heterogeneity across individuals and (b) a reference point that is *partially* sticky for most individuals. This heterogeneity, specifically in loss aversion, implies that a significant proportion of individuals have moderate loss aversion and hence invest in stocks even for small risk premiums. The partially sticky reference point implies that stock payoffs straddle the reference point even after high gains, so loss aversion compounds risk aversion and individuals exhibit the disposition effect even for high risk premiums.

This paper is closely related to two strands of the empirical literature: studies that structurally estimate prospect theory preferences and studies on the dynamics of risk-taking. Most structural estimation studies use data from lab experiments or TV game shows (e.g., Hey and Orme 1994; von Gaudecker, van Soest, and Wengstrom 2011; Post et al. 2008). At the cost of losing experimental control but with the benefit of capturing natural behavior, a small number of studies estimate prospect theory preferences in the field. Most of these use price data to elicit the preferences of a representative agent: Jullien and Salanie (2000) and Snowberg and Wolfers (2010) use horse race track betting prices, and Kliger and Levy (2009) and Polkovnichenko and Zhao (2013) use option prices. An exception is Barseghyan et al. (2013), who use individuals’ insurance choices to estimate heterogeneous risk preferences. Our paper complements these studies, and has two distinct features. First, we observe a large number of choices per individual, enabling us to estimate for the first time how preferences not only vary across individuals but also over time. Second, we observe a great variety of lotteries in terms of payoffs (both positive and negative) and probability distributions, enabling us to estimate all features of prospect theory for a wide range of prizes and probabilities. Using the same data, Andrikogiannopoulou and Papakonstantinou (2016) develop a model with more comprehensive preference heterogeneity, but no history dependence, to assess the relative prevalence of prospect theory’s main behavioral features: loss aversion and probability weighting. A key element of their analysis is the development of a mixture model of behavioral types and the allocation of individuals to these types. They find that, while loss aversion is

important, probability weighting is much more prevalent. In contrast, here we use CPT as the starting point, and we fully exploit the panel structure of the data to estimate history-dependent preferences.

The second strand of related literature has documented that past events significantly affect subsequent risk-taking. In an experiment, Thaler and Johnson (1990) have shown that people take more risk after gains and also when they have a chance to recover prior losses; similar behaviors have been observed in TV game shows (e.g., Gertner 1993; Post et al. 2008). In the field, various studies have documented history dependence. Calvet, Campbell, and Sodini (2009) find that prior gains (losses) increase (decrease) the risky share in households' portfolios. Kaustia and Knupfer (2008) and Malmendier and Nagel (2011) find that individuals who experience positive (negative) returns in initial public offering (IPO) and stock investments, respectively, exhibit higher (lower) future propensity to take related risks. 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) find the opposite. Notably, these studies estimate an average relationship between prior outcomes and risk-taking using reduced-form regressions.⁵ In contrast, our structural approach allows us to estimate the underlying preferences that give rise to this behavior.

Our paper also contributes to the literature on the characterization of the reference point. A variety of asset-pricing models with dynamic referents have been proposed in the finance literature. One class of models assumes that the referent is forward-looking and depends on the expectation of future outcomes as proposed by Kőszegi and Rabin (2006, 2007, 2009); see Andries (2012) and Pagel (2016). The evidence on such forward-looking reference points is mixed (see O'Donoghue and Sprenger 2018 for a review), while their plausibility in finance is unclear, because, as Barberis (2018) notes, the high degree of uncertainty in the stock market may hamper their calculation and use. Instead, most dynamic implementations of prospect theory in finance use a backward-looking reference point, consistent with the aforementioned evidence that past events affect subsequent risk-taking. For example, Barberis, Huang, and Santos (2001) suggest that an investor's gain/loss utility from asset returns depends on a measure of his historical investment performance that adapts sluggishly to past gains/losses. In this paper, we find evidence from a market setting that individuals' behavior depends on a

⁵ An exception is Post et al. (2008), who estimate CPT with no heterogeneity, no probability weighting, and a reference point/loss aversion related to future expectations. Their implementation of history dependence is also specific to their game show setting.

backward-looking, adaptive reference point that is most similar to that of Barberis, Huang, and Santos (2001).

We note that the preference estimation literature, 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 buying insurance), so the derived estimates may not be generalizable to the entire population. Indeed, wagering markets may attract individuals with lower loss aversion and/or higher probability weighting, which make them more amenable to returns with negative mean and positive skewness. However, the popularity of sports wagering, together with the fact that our estimates are not only heterogeneous but also within the range of experimental estimates, suggest that a potential selection bias is likely not severe.⁶ The second potential criticism is that each study uses data from a specific setting, so the derived estimates may not be generalizable to other domains. For example, in the insurance context, individuals may be primed to make choices that limit their exposure, so estimates may be less relevant in, for example, the stock market where people's primary motive is financial gain. Similarly, one could argue that one of the drivers of sports wagering behavior is nonpecuniary, for example, entertainment. While this means that estimates from sports wagering—or any other context—may not be used universally, it also means that they provide a valuable complement to the existing estimates in the literature.

Importantly, our preference parameter estimates may 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, and sports handicappers play the role of financial analysts. Second, individuals who wager on sports exhibit characteristics similar to those who participate in traditional financial markets. For example, individuals in our sample have similar demographics to individuals who trade stocks online: they are likely to be male and younger than the general population. Furthermore, survey evidence shows that, like in the stock market, the primary driver of behavior in the sports wagering market is financial gain (“to make money” or to “[win] big money”) and nonpecuniary motives, such as entertainment and team loyalty, are important but secondary.⁷ Conversely, recent studies show that many individuals in the stock market also view trading as an

⁶ A 2008 survey in the United States (by ESPN) shows that 50% of adults wager on sports each year, and it is estimated that about \$1 trillion is wagered on sports globally per year (see the H2 Gambling Capital 2013 report).

⁷ For survey evidence on sports wagering behavior, see the 2010 British Gambling Prevalence Survey. Consistent with this, we find that team loyalty/fandom does not play a strong

entertaining activity, are motivated by loyalty, and prefer stocks with lottery characteristics.⁸ Anecdotal evidence also suggests that many stock traders engage in gambling activities (see McDonald and Robinson 2009). Third, at first glance, it seems that participants in a market with negative expected returns exhibit risk-seeking behavior which is at odds with the major stylized fact of a positive premium in the stock market, but behavior is actually consistent across the two markets. Indeed, we estimate a generally concave value function and substantial probability weighting, implying that behavior in our market is better explained by a preference for skewness, not risk. Barberis and Huang (2008) show that these preference characteristics—aversion to risk coupled with preference for skewness—imply not only a positive equity risk premium but also a negative premium for idiosyncratic skewness, which is consistent with the low mean returns of, for example, IPOs, distressed stocks, and individual stock options.

1. 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.

1.1 The sports wagering market

The sports wagering market is a quote-driven dealer market 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, that is, 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.⁹ 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

role in our data, since, on average, individuals place fewer than 3% of their wagers on any one team.

⁸ In a survey, Hoffman (2007) finds that the second strongest motive for investing is that it is “a nice free-time activity,” behind “financial gain” but ahead of “safeguard[ing] retirement.” Dorn and Sengmueller (2009) find that most investors enjoy investing and that entertainment partially drives trading. Gao Bakshi and Lin (2015) find that, when there is a large jackpot, some individuals substitute buying lottery tickets for trading stocks. Kumar (2009) shows that individuals prefer stocks with high idiosyncratic volatility and skewness.

⁹ This describes fixed-odds betting, which is commonly used in sports wagering. In contrast, in parimutuel betting, which is more commonly used in horse race track betting, individuals place bets by putting money in a pool, and payoff odds are determined when the pool is closed, and odds are calculated by sharing the pool among the winners.

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, an individual clearly 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 associated with perfectly balancing the book at all times, while it is not particularly risky for the bookmaker because the commission he charges provides a cushion against the unbalanced liabilities implied by an unbalanced book. Overall, it could lead to greater long-run profits.¹⁰ 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 3.5, our results are robust to alternative behaviors.

1.2 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

¹⁰ Other 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.

(October 2005 – November 2010).¹¹ 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 shown in Figure 1, which summarizes all the relevant information.

[Figure 1 about here]

For each bet placed by an individual in our sample, we observe the following: (a) bet date, (b) bet event and outcome chosen, (c) bet amount, (d) bet type, (e) prices for all outcomes of the bet event, and (f) bet result. In addition, we have information about the gender; age; and country and ZIP code of residence of the individuals. Table 1 presents 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 similar to those reported for samples of individuals who invest in the stock market through online brokers (e.g., Barber and Odean 2002; Glaser 2003).¹² An individual, on average, has wagered on 5 events during the same day, has wagered in the sportsbook on 35 different days, and has a betting frequency of approximately once a week.

[Table 1 about here]

1.3 Lottery representation

In this section, we explain how we represent individual choices as lotteries, that is, 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. Like in any field setting, it is not possible to know individuals' subjective beliefs and hence the probabilities individuals associate with these prizes. Because observed behavior

¹¹ The original data set contains 400 individuals, but we drop 64 individuals, who placed fewer than 5 wagers.

¹² In 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.

can be explained by several combinations of preferences and 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 race track betting, insurance, or options trading—resolve this identification problem by assuming beliefs are homogeneous and coincide with the (econometrician’s approximation to the) rational beliefs (see Barseghyan et al. 2018 for a review).¹³ In our main analysis, we make use of a key advantage of the market we study, which is that prices are quite efficient so the true probabilities of the events can be readily approximated by the probabilities implied by their betting odds (see Section 3.5 for details).¹⁴ Furthermore, the following observations suggest that individuals’ subjective probabilities should not deviate significantly from the implied probabilities. First, the betting platform provides access to calculators that convert quoted prices into their implied probabilities, so it is likely that individuals observe these probabilities before placing their wagers. Second, we find that individuals spread their wagers across various leagues and teams: on average, an individual in our sample has wagered on 34 (182) different leagues (teams), while less than 20% (3%) of his wagers are on any specific league (team), making it unlikely that individuals have superior information, either real or perceived, about any specific league and/or team. In Section 3.5, we show that our results are robust to (1) allowing for small but significant deviations of the subjective from the implied probabilities and (2) approximating the true probabilities with the win frequencies of past outcomes with similar prices.¹⁵

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

¹³ This 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.

¹⁴ For comparison, studies in the insurance setting typically assume that an individual’s belief about his claim rate equals the true probability, which is estimated from 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 a counterparty’s behavior, Post et al. (2008) propose a rule that explains 74% of the variance of the counterparty’s behavior and assume that contestants’ beliefs are consistent with this rule.

¹⁵ The latter approach is employed by Jullien and Salanie (2000) and Snowberg and Wolfers (2010) in studies of parimutuel betting in the horse race track, a setting in which the market is quite inefficient and hence implied probabilities significantly deviate from the true probabilities. Prices are quite efficient in our fixed-odds betting market, so it is not surprising that our results are robust.

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 that is defined as the set of all bets placed over a single day by an individual.¹⁶ The lottery representation is constructed using all possible combinations of payoffs from all the bets placed during the session. After dropping 1,043 lotteries that are difficult to compute, we have a final sample of 11,490 lotteries chosen by 336 individuals.¹⁷ Table 2 reports 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 a negative mean because the bookmaker's commission is embedded in the odds. 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.

[Table 2 about here]

2. 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

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

¹⁶ Defining the play session as containing only individual bets would be too narrow, because multiple bets are often placed simultaneously. At the other extreme, defining the play session as containing all bets that are not yet settled is similar to our definition of the play session, because most bets are placed on the day of the event.

¹⁷ We 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.

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} are observables that include linear and quadratic terms of the cumulative gains and losses realized over a look-back period of t calendar days prior to play day l , a_n denotes individual fixed effects, and b and ζ_{nl} are the regression coefficients and errors. We consider look-back periods with t equal to 7, 30, and 365 calendar days; all yield qualitatively similar results. 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. 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 following 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.

Table 3 presents results from these two models, with previous gains and losses measured over a look-back period of 7 calendar days. 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 and 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), which, because of the bookmaker's commission, leads to a commensurate reduction in the mean payoff of their selected lottery, and individuals 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.

[Table 3 about here]

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.

3. 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.

3.1 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.

To specify the set of day lotteries from which individuals make their choices, we assume that individuals may consider any lottery that can be constructed in the sportsbook, but they are less likely to consider

lotteries that are rarely chosen by any individual on any day in our sample, for example, lotteries that involve a large number of complex bet types.¹⁸ Hence, we generate a set of lotteries by randomly drawing lotteries' risk characteristics (i.e., the number and types of bets they involve, the odds of each wager, and the amount staked) from the empirical distributions of the characteristics of the lotteries chosen by individuals in our data.¹⁹ Then we augment this set with the chosen lottery and the safe lottery, which represents the option not to place a bet. In constructing this choice set, we consider solely the lotteries' risk characteristics rather than the underlying match characteristics (e.g., the specific teams wagered) because, as discussed in Section 1.3 above, individuals do not seem to have a systematic preference toward specific teams/events but rather combine wagers on various teams and events. As a result, the choice set does not change over time, because the same set of payoff risk profiles can be constructed at any time in the sportsbook by combining the multitude of available prices with different stakes and under different bet types.

3.1.1 Preferences. Next, we specify individuals' preferences over lotteries. Motivated by our 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 cumulative prospect theory (CPT), proposed by Tversky and Kahneman (1992) (henceforth TK). Specifically, at time t , an individual with wealth W_t and preference parameter vector $\theta := (\alpha, \lambda, \gamma, \kappa, \delta)'$ evaluates lottery L with probabilities $\{p_i\}$ and payoffs $\{z_i\}$ with $z_{-m} \leq \dots \leq z_0 = 0 \leq \dots \leq z_M$ as follows. He computes the difference between his final wealth $z_i + W_t$ and a reference point W_t^{RP} , which depends on the rate of memory decay δ and state variable s_t , and assigns 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 and

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

¹⁸ Even 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.

¹⁹ See Appendix B for details on the exact procedure and distributions used to generate the choice set. Also, this procedure is essentially a reduced-form implementation of a structural model of choice set generation in which the likelihood of a lottery being included in the choice set decreases with its complexity; 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. The two approaches would yield similar choice sets, so we have elected to use the reduced-form implementation to reduce computational complexity.

is the decision weight for payoff z_i , with $w(\cdot; \gamma, \kappa)$ the probability weighting function. As an aside, in this subsection we simplify notation by omitting the subscript n denoting the individual.

For the value function, we use the form proposed by TK:

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

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, that is, 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 deviations in cumulative probability from the natural boundaries of certainty and impossibility, and κ captures the sensitivity to extreme outcomes. Figure 2 shows 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

²⁰ In the original CPT formulation, λ is constrained to be greater than one (hence termed loss *aversion*) to reflect the observation that “losses loom larger than gains” (e.g., the endowment effect). We do not place this restriction on λ , because recent experiments (e.g., Abdellaoui, Bleichrodt, and Paraschiv 2007; von Gaudecker, van Soest, and Wengstrom 2011) have found that a small but significant proportion of individuals exhibit gain seekingness, that is, $\lambda < 1$. As we discuss in Section 3.3, we estimate that only a small proportion (about 10%) of individuals have $\lambda < 1$; reestimating the model with the constraint $\lambda \geq 1$ has no significant effect on our results.

concave which corresponds to an overweighting of the right (left) tail of the distribution of positive (negative) outcomes; as κ increases, the concavity of the function increases. For $\kappa < 1$ (not shown), the function becomes globally convex.²¹

[Figure 2 about here]

3.1.2 Reference point. The original CPT pertains to static choice settings and proposes that the reference point that separates losses from gains could equal the status quo, that is, the individual's current wealth ($W^{\text{RP}} := W$). Extending prospect theory to dynamic settings, a variety of models for the reference point have been proposed in the finance literature. On the one hand, there are forward-looking models in which the reference point relates to the expectation of future outcomes as proposed by Kőszegi and Rabin (2006, 2007, 2009). For example, in Andries (2012) investors care about actual relative to expected consumption while in Pagel (2016) they derive gain/loss utility by comparing beliefs before and after receiving some news. These models successfully match the observed moments of stock returns, though the former does not generate return predictability and the latter has some difficulty generating a low and stable risk-free return. As we noted in the introduction, such forward-looking models have found mixed success in explaining observed behaviors in different settings, while their plausibility in finance is unclear given the high degree of uncertainty in the stock market. On the other hand, motivated by experimental evidence that prior outcomes affect individuals' subsequent risk-taking (Thaler and Johnson 1990; Gertner 1993), most implementations of prospect theory in finance use a backward-looking reference point. For example, Barberis, Huang, and Santos (2001) suggest that the gain/loss utility that an investor derives from asset returns depends on a measure of his historical investment performance which is assumed to adapt sluggishly to his past gains/losses and captures his memory thereof. In calibrations, this model can match the high mean, volatility, and predictability of stock returns, while maintaining a low and stable risk-free return.²² In this paper, we model a backward-looking, adaptive reference point similar to that of Barberis, Huang, and Santos (2001),

²¹ To reduce the number of parameters and facilitate identification, we deviate from the original CPT formulation in which the value and probability weighting functions have different parameters for gains/losses. Indeed, TK estimate the gain-/loss-specific parameters to be similar, so many studies that estimate or use CPT preferences make similar simplifying assumptions.

²² Outside the prospect theory paradigm, Campbell and Cochrane (1999) use a backward-looking referent/habit that depends on the history of aggregate consumption, and Routledge and Zin (2010) use as referent the certainty equivalent rather than some historical or expected future outcome. In calibrations, both models match the basic

and we estimate an individual-specific parameter capturing how far back the investor recalls his past gains/losses and therefore the adaptation speed of his reference point.

Specifically, 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 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.²³

3.1.3 Identification. 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, that is, 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

moments of stock returns, but they also yield a counterfactually high correlation between stock returns and consumption growth.

²³ Though 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 and hence behaviors like the house money and break-even effects (see Section 3.6).

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.

3.2 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. Specifically, 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.

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 $\theta_n := (\alpha_n, \lambda_n, \gamma_n, \kappa_n, \delta_n)'$ are the preference parameters 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 be interpreted as an econometric error necessary to reconcile the data with a theory of choice. As is standard in the discrete choice literature, we assume 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 individual n chooses lottery L_j on date t out of choice set 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, for example, to 200, the number of clusters has a very small effect on our results; see also Appendix B 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. Specifically, because 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.²⁴

As already mentioned, to study behavior we use information in individuals' observed choices, that is, both the wagers they place and the frequency with which they place wagers. Specifically, 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 can bet on date t , he either (a) chooses the safe lottery L_0 that pays 0 with

²⁴ Von 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 and gender).

certainty, and, hence, we observe no bet on date t , that is, $x_{nt}=0$, or (b) he chooses lottery $y_{nt} \neq L_0$, and, hence, we observe a bet on date t , that is, $x_{nt}=1$. If he does not have the opportunity to bet, we observe no bet on date t , that is, $x_{nt}=0$. Thus, the likelihood of observing (x_{nt}, y_{nt}) given the individual's parameters and state is

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

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, that is,

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

Our baseline priors are $\underline{\kappa} = \mathbf{0}$, $\underline{K} = 100I$, $\underline{\lambda} = \text{rank}(\Sigma_{\tilde{\vartheta}})$, 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 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 relegate a detailed description of our algorithm to Appendix A. Next, we present our results.

3.3 Estimation results

Panel A of Table 4 presents the posterior estimates for the population means of the model parameters, and Table 5 shows the medians of the estimated population distributions. Our population estimates for the value function 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). Our 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. Figure 3 plots the value and probability weighting functions corresponding to various combinations of α , λ , γ , and κ .

[Table 4 about here]

[Table 5 about here]

[Figure 3 about here]

Using hypothetical lotteries, TK 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 the probability weighting function has a more pronounced inverse S shape so additionally overweights outcomes near the reference point. Subsequent experiments have produced a wide range of CPT parameter estimates: According to a meta-analysis 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 λ in the range [1.07,3.2].²⁵ For the probability weighting function, different studies have used different functional forms; estimated parameters usually indicate an inverse S shape, while a few studies do not find support for a change in concavity (e.g., Abdellaoui, Vossman, and Weber 2005). Thus, though some of our estimates deviate from the widely used TK estimates, they are well within the range of estimates found in the experimental literature. We conclude then that our findings validate the general experimental finding that individuals are risk averse (loving) over gains (losses), exhibit loss aversion, and overweight the probabilities of extreme outcomes, and we extend support for this finding to a real market setting.

We do note that our estimate of loss aversion lies toward the lower end of the range of existing estimates. A possible explanation is that existing estimates of loss aversion are predominantly from experiments rather than from a real-world market setting as the one we observe. Another possibility is that more loss-averse people may not participate in a sports wagering market. Yet it is reassuring that our estimated loss

²⁵ Specifically, 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]; and 3 studies estimate an average λ below 1.5 and 3 estimate it above 2.5.

aversion is quite heterogeneous across individuals, and also that it is only weakly related to individuals' propensity to accept betting opportunities and it is unrelated to the propensity to wager on popular matches. These findings indicate that a potential sample selection is unlikely to be severe.²⁶ Furthermore, as discussed in the introduction, the similarities between sports wagering markets and traditional financial markets indicate that people who participate in these markets are likely self-selected in a similar manner. Therefore, our estimates may be more relevant for financial markets than corresponding estimates from other settings (e.g., lottery choices in experiments or deductible choices in insurance).

Regarding the memory decay parameter δ , we estimate its population mean/median to be 0.60/0.66, meaning most individuals slowly absorb past profits and the reference point they use to separate gains from losses tends to stick to past levels. For example, $\delta=0.66$ implies that the reference point absorbs 34%/56%/87% of the profits realized 1/2/5 play days earlier. That is, after a gain of, for example, €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 treated 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 that incorporate memory in behavioral finance have arbitrarily assumed individuals have a perfectly sticky reference point, that is, $\delta=1$. But as we show in Section 4 on the disposition effect, the assumption that $\delta=1$ may lead to qualitatively different conclusions than parameters informed by actual estimates. Furthermore, $\delta=1$ cannot be thought of as corresponding to one extreme in behavior (with $\delta=0$ corresponding to the other extreme), because δ may affect behavior nonmonotonically.

Panel B of Table 4 presents posterior estimates for the model parameters' population variances. Table 5 presents percentiles of the model parameters' estimated distributions. Figure 4 plots their estimated densities. We find substantial heterogeneity across individuals in all preference parameters, which is mostly ignored by applied models in finance. Specifically, our estimates imply that the 25th/75th percentiles of the value function parameters are 0.83/0.90 for α and

²⁶ For example, individuals in the top 5% of propensity to accept betting opportunities (measured as observed betting frequency over estimated probability of having a betting opportunity) have average λ of 1.37, while those in the bottom 5% have average λ of 1.68. Though it is expected that loss aversion affects participation, comparing this variation with the heterogeneity in λ across our sample (5th percentile of 0.92, 95th percentile of 2.35), we see that the latter is much wider, indicating that the participation decision is not predominantly driven by loss aversion. Similarly, it is reassuring that we find no relation between our preference estimates and the propensity to place wagers involving popular teams, as it could be theorized that individuals placing such wagers are likely to be more representative of the wider population.

1.18/1.70 for λ , and for the probability weighting function parameters they are 0.85/0.94 for γ and 1.00/1.59 for κ . In Figure 3, we present plots of the value and probability weighting functions corresponding to these percentiles. Regarding the memory decay parameter δ , we find 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. Specifically, 15% of our individuals have very short memory ($\delta < 0.1$), 22% have very long memory ($\delta > 0.9$), and the remaining 63% have memory that decays at widely different 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 our choice of priors, because in a prior sensitivity analysis we find that the priors have a small effect on our posteriors. Figure 4 also shows 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 we use in our analysis and hence is the likeliest to affect the posteriors; we find that the estimated 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 ($\underline{\Lambda} = 10I$).

[Figure 4 about here]

3.4 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 given 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 or 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.

Panel A of Table 4 shows 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, for example, 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 (e.g., von Gaudecker, van Soest, and Wengstrom 2011). Nonetheless, this small random 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—that is, 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.²⁷ Furthermore, Figure 5 plots the histogram of the mean ratio ranking of the chosen lotteries across all choices for an individual. 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%.

[Figure 5 about here]

We also compare the fit of our model with that of alternative models: (a) one in which lotteries are chosen randomly; (b) one in which individuals are risk neutral, (c) one in which preferences are homogeneous, and (d) one in which preferences are heterogeneous, but not history dependent. Table 6 shows 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, and, hence, it also selects our model, which incorporates both heterogeneity and history dependence.

[Table 6 about here]

²⁷ To 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.

3.5 Subjective beliefs

As discussed in Section 1.3, observed 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 by using the standard assumption of rational beliefs, which we approximate with the match outcome probabilities implied by the quoted prices. Here, we first show that the market under study is quite efficient, so the implied probabilities are indeed close to the true. Then we show our results are robust to (1) approximating outcomes' true probabilities with the win frequencies of past outcomes with similar prices and (2) 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.²⁸ We divide all possible outcomes of these matches into 100 percentile odds groups. Figure 6 plots, for each of these groups, the mean realized return in excess of the mean 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, that is, the return from betting on short odds is higher than that from betting on long odds, this is not very pronounced. Specifically, a number of longshot percentiles have the same estimated excess return as the favorite percentiles, whereas the 95% confidence intervals show that the excess return for most percentiles is not significantly different from zero. Indeed, the FL inefficiency we find is much less pronounced than that found in studies that use data on parimutuel betting in horse races (see, for comparison, Figure 6, in which we superimpose the corresponding plot for the Snowberg and Wolfers 2010 data on parimutuel betting).²⁹ Nonetheless, to examine whether the mild FL inefficiency present in our market affects our results, we repeat our analysis replacing the implied probabilities with the win frequencies of

²⁸ We obtained the data from the following sources: (1) football-data.co.uk, (2) matchstatistics.com, and (3) 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.

²⁹ Although the literature has consistently found large pricing inefficiencies in parimutuel horse race track 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.

past outcomes with similar prices. Panel A of Figure 7 shows that our results remain largely unchanged under this alternative assumption.

[Figure 6 about here]

[Figure 7 about here]

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. Specifically, letting $1-p$ be the implied win probability for an individual's selected outcome, we let the "perceived" win probability of this outcome to equal $1-\zeta p$ for $\zeta < 1$. Panel B of Figure 7 plots 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 values Barcelona more highly than the market), 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), with each wager involving a different level of risk. Furthermore, individuals typically place multiple wagers on the same day, so they also need to decide if and how they will combine these under different bet types involving different levels of risk. Thus, even if our assumption of rational beliefs is not exactly correct, the chosen day lotteries are still informative about individuals' risk preferences.

3.6 Intuition and examples

Here, we present intuition about the effects of the CPT preference parameters and of past profits on individuals' subsequent risk-taking. We focus on three individuals from our sample for whom we have estimated different value function curvature, loss aversion, and memory decay parameters. We analyze their behavior in two hypothetical states in which they face the option of accepting or rejecting the same binary lottery; in one state, individuals have no prior 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 lottery's payoffs.

The first 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 past wealth ($\delta \approx 1$). The second individual (*B*) has a similar value function but his reference point partially updates to incorporate past outcomes ($\delta=0.50$). The third

individual (C) has a perfectly sticky reference point, like individual A , but a highly curved value function ($\alpha=0.50$) and low loss aversion ($\lambda=1.12$). To facilitate the exposition, we select these three individuals such that none of them distorts probabilities; that is, γ and κ are close to 1. The binary lottery we consider pays -7 or 10 , with equal probability, chosen such that all three individuals are indifferent between choosing it or rejecting it in the neutral state, which facilitates our comparison of risk-taking across states.

Panels A, C, and D of Figure 8 plot an individual's value function and the lottery's certainty equivalent in the neutral state and after a large prior gain of 8. In the gain state, individuals A and C do not integrate the gain to their reference point, and, hence, they perceive the lottery's outcomes as gains ($8-7=1$ or $8+10=18$), so loss aversion becomes irrelevant. If the value function's concavity is mild relative to loss aversion, as is the case for individual A , the lottery appears more appealing than in the neutral state; if it is strong, as is the case for individual C , the lottery appears less attractive. Thus, individual A exhibits increased risk-taking after gains, consistent with the house money effect, whereas individual C exhibits the opposite behavior. In contrast, individual B integrates half of the gain into his reference point, so he perceives the lottery's outcomes as $4-7=-3$ or $4+10=14$. Because these straddle the reference point, loss aversion compounds the value function curvature, so for individual B the lottery appears less appealing than in the neutral state.

[Figure 8 about here]

Panel B of Figure 8 plots individual A 's value function and the lottery's certainty equivalent in the neutral state and after a small gain of 4. While A accepts the lottery after a large gain, he rejects it after a small one. Comparing this behavior with that of individual B after a large gain, we see that a smaller prior gain has the same effect as a less sticky reference point, as in both cases the (perceived) payoffs straddle the reference point, so the interplay of loss aversion and risk aversion makes the lottery less appealing than in the neutral state.

A similar reasoning can be used to analyze risk-taking behavior after losses. For example, after a large loss, individual A would exhibit increased risk-taking, consistent with the break-even effect, as the convexity of the value function over losses dominates his decision, while the opposite would be observed after a small loss.

These examples show that, in the presence of heterogeneity in risk preferences and in the level of prior profits, estimating an average reduced-form relationship between past profits and risk-taking can be misleading. This highlights the importance of structural estimation of

heterogeneous history-dependent preferences, which can help us better understand and predict behavior across a variety of situations. Next, we use our estimates to study individuals' dynamic trading behavior.

4. 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.³⁰ The results of these studies are largely based on (a) the assumption that all investors have the same preferences, (b) the median CPT parameter estimates from the TK experiment, and (c) (in dynamic settings) the assumption that individuals have either no memory or perfect memory. In this section, we use the individual-level preference parameters we estimated in Section 3 to study individuals' dynamic trading behavior, with a particular focus on the disposition effect, that is, the tendency of individuals to sell (retain) stocks whose value has increased (decreased) since purchase. Our goal is to study how the preference heterogeneity we estimate translates to heterogeneity in exhibiting the disposition effect, how the estimated history dependence affects the disposition effect, and more generally whether, with these estimates, prospect theory can explain the prevalence of this behavior.

The disposition effect (henceforth DE) is a robust empirical finding regarding trading behavior.³¹ 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, that is, the purchase price. Because he is risk averse (risk seeking) over gains (losses), he prefers to sell (retain) the stock, so he exhibits the DE. However, this argument ignores how loss aversion interacts with assets' risk characteristics to affect investors' decisions. Barberis and Xiong (2009) use a simple model to show that, contrary to conventional wisdom, an investor with CPT preferences defined over annual trading profits, preference parameters as estimated by TK, and a perfectly sticky reference point would exhibit the opposite behavior to the DE. That is, he would optimally increase (decrease) his holdings of a stock after it appreciates (depreciates). The

³⁰ See, for example, Polkovnichenko (2005) and Mitton and Vorkink (2007) for applications of prospect theory in stock market participation and under-diversification; Kyle, Ou-Yang, and Xiong (2006) and Li and Yang (2013) for applications in dynamic trading behavior; and Barberis and Huang (2008) and Zhang (2006) for applications in the equity premium puzzle and the pricing of idiosyncratic skewness.

³¹ See 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 mutual fund managers; and Weber and Camerer (1998) for experimental evidence.

intuition behind this is the following. When the investor first decides whether or not to buy a stock, he faces a lottery whose payoffs straddle the reference point. Hence, loss aversion “bites” and may 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 retaining 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 panel A of Figure 8, which depicts the choice to take risk after a “large” gain.

4.1 Model setup

Following Barberis and Xiong (2009) (henceforth BX), we consider a discrete-time binomial model with dates $t \in \{0, 1, \dots, T\}$, with 0 the beginning and T the end of the year. At each $t \in \{0, \dots, T-1\}$, there is a risk-free asset with (gross) return R_f and a risky asset with (gross) return $R_t \in \{R_d, R_u\}$ such that its excess return has annualized mean μ and standard deviation σ , with $0 < R_d < R_f < R_u$. To maintain comparability with the results in BX, we set R_f to 1, μ in the range 3% to 13%, and 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 0.5.³² Furthermore, we set $T=2$, that is, 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.³³

At $t=0$, the unit price of the risky asset is (an arbitrary) $P_0 > 0$ and the investor has initial wealth (an arbitrary) $W_0 > 0$.³⁴ At dates $t \in \{0, 1\}$ the investor chooses the proportion of wealth χ_t that he invests in the

³² Using CRSP data on the risk-free (the 1-month Treasury-bill) rate and on stock returns 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%. Calibrating the model to these values, our results below are even stronger.

³³ In Barber and Odean (2000), the mean/median U.S. household holds 4.3/2.6 stocks and trades 5 times per year, corresponding to 1 or 2 trades per stock per year. In Alvarez, Guiso, and Lippi (2012), the median number of yearly trades for Italian equity investors is 2.

³⁴ Initial 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, and, hence, optimal actions do not depend on W_0 .

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

$$W_{t'} = W_{t'-1} [1 + \chi_{t'-1} (R_{t'} - 1)], \quad \text{for } t' > t \quad (\text{budget constraint})$$

$$W_T \geq 0 \quad (\text{no bankruptcy})$$

$$\chi_t \geq 0 \quad (\text{no shorting})$$

where U is the CPT utility functional defined in Section 3.1, that is, $U(\theta, L_t^T, s_t) = \sum_i w_i v(W_T - W_t^{\text{RP}})$ 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) to (5).³⁵ Specifically, at $t=0$, the reference wealth equals contemporaneous wealth, and, at $t=1$, it equals a convex combination of contemporaneous wealth and wealth at $t=0$, that is,

$$W_0^{\text{RP}} := W_0$$

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

where δ is the memory decay parameter introduced in Equation (5). We note that probability weighting ($\gamma \neq 1$ and/or $\kappa \neq 1$) and/or a reference point that is not perfectly sticky ($\delta < 1$) generate a time inconsistency: the optimal investment level at $t=1$ from the $t=0$ perspective is different from that from the $t=1$ perspective. We focus on the behavior of agents who are sophisticated in that they can foresee this time inconsistency and who—as is usually the case—have no access to a commitment device. In the Internet Appendix, we also present results for agents who are either naive or sophisticated and *can* commit to an investment plan (for a similar analysis, also see Barberis 2012); qualitatively, the results are similar across all cases; however, we discuss some quantitative differences in the Internet Appendix.

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 ($R_1 =$

³⁵ Like the investor in BX, ours also can borrow money. But because, 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, so we also present results for the case with a no-borrowing constraint, that is, $\chi_t \leq 1$. BX formulate the problem and define the DE by the number of shares bought rather than the proportion of wealth invested in the risky asset; the two definitions yield qualitatively similar results for the case with borrowing, and identical results for the case without borrowing.

R_u) state at $t=1$, and then the optimal action χ_0^* at $t=0$. Finding these 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 an agent who (a) at $t=1$ uses a reference point that perfectly sticks to the original wealth level at $t=0$, which corresponds to $\delta=1$; (b) has a value function with the TK median parameter estimates, that is, $\alpha=0.88$ and $\lambda=2.25$; and (c) does not distort probabilities, which corresponds to $\gamma=\kappa=1$.

Before we present our results, we note the following difference between the model of behavior in this application and in the estimation (see Section 3.1). 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. 2018) it seems most sensible that individuals narrowly frame each day’s choices among 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 TK parameter estimates which are similarly derived from choice situations involving narrowly framed short-lived lotteries.

4.2 Results

Table 7 presents, 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 retaining 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 like

those in BX, the magnitudes and hence our conclusions, are different. Specifically, BX find that the 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 Comparing the results in the two panels of Table 7 we see that, for most values of μ , participation is the same with and without borrowing, but the prevalence of the DE is about 20% lower, though still substantial, for the case with borrowing. This is because the possibility of borrowing makes it more likely that an individual is mechanically forced to reduce his stock holdings after a loss to satisfy the no-bankruptcy constraint. As a result, though BX raise valid concerns about the ability of prospect theory to explain the DE, our message is more encouraging, because based on our estimates prospect theory *can* explain the DE for a much wider range of stock return parameter values.

[Table 7 about here]

Next, we examine why our results differ from those in BX. As mentioned above, the 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 μ s. Therefore, we focus on why individuals with our estimated parameters participate for small μ and exhibit the DE for both small and large μ . Also, while 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, we estimate a significant proportion of individuals with loss aversion λ below the value 2.25 used by BX.³⁸ These individuals buy the

³⁶ In addition to their main specification, BX also propose a model similar to prospect theory but in which utility is defined over realized gains and losses. The authors 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 $\mu > 11\%$.

³⁷ Little empirical evidence addresses the DE's prevalence in the population. Using a subsample of the Barber and Odean (2000) data, Dhar and Zhu (2006) calculate that 80% of the individuals exhibit the DE. Considering that this proportion is likely to be upward biased (as the authors note), because of the systematic exclusion of individuals from the original sample, it is not far from our estimated proportion of about 60% for plausible equity premium and Sharpe ratio parameters.

³⁸ A low average λ is not necessary for obtaining a high participation rate. Rather it is sufficient that there is substantial heterogeneity in λ such that a significant proportion

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. Panel A of Figure 9 plots the proportion of individuals who optimally participate in the stock market at $t=0$ as a function of λ and μ , fixing the other parameters to our individual-specific estimates. We see, for example, that λ s below 2.0 are sufficient to generate some participation for reasonable μ s, and λ s below 1.7 are sufficient to generate some participation for all μ s. Thus, allowing for heterogeneity in λ is mostly responsible for the higher participation rates, so for our finding that prospect theory can explain the DE for low values of μ .

[Figure 9 about here]

Second, we estimate that many individuals have memory decay parameter δ significantly below the value 1 used by BX. Panel B of Figure 9 plots the proportion of individuals who optimally exhibit the DE (conditional on participating) as a function of δ and μ , fixing the other parameters to our individual-specific estimates. We observe an inverse-U relationship between δ and the proportion exhibiting the DE. On one end, $\delta=0$ implies no DE as the optimal decision after gains and losses is the same; on the other end, $\delta=1$ implies no DE for large μ s, consistent with BX's insight described earlier. But 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 compounds risk aversion, making the value function even more concave (for an illustration of this intuition, see panels A and C of Figure 8). For example, for $\mu=10\%$ (the lowest risk premium with participation, but no DE, with the BX parameters) we find that $\delta=0.7$ yields a DE rate of 66%, while $\delta=1$ yields a DE rate of 6%. Thus, the intermediate values of the memory decay parameter δ that we have estimated for most individuals are mostly responsible for our finding that prospect theory can explain the DE for high values of μ .

5. Conclusion

In this paper, we develop a structural model of dynamic choice under risk within the prospect theory paradigm, and estimate it using data on individual trading activity in the sports wagering market. We find that individuals in this market exhibit typical prospect theory features: they are averse to risk, they have a strong preference for skewness, and they evaluate gains/losses relative to a reference point that depends

of individuals has $\lambda < 2.25$ and optimally buys the stock. Indeed, all experiments studied by Booiij, Van Praag, and Van De Kuilen (2010) find large heterogeneity in λ , some even finding significant proportions of individuals with λ below 1.

on their past performance. The parameters we estimate are similar to those in the experimental literature and consistent with several stylized facts about financial markets. Hence, the evidence in this paper points to a unified prospect theory-based framework that can be helpful for understanding people's risk-taking across a variety of domains, that is, laboratory experiments, sports betting, and financial markets.

Our results also suggest an important role for preference heterogeneity across individuals as well as for history dependence in individual behavior. Importantly, we estimate the distribution of the degree of stickiness of the reference point individuals use to separate gains from losses. We find that, for most individuals, the reference point is somewhat, but not perfectly, sticky. By characterizing the dynamics of the reference point, our paper provides field evidence on a parameter that is of central importance in prospect theory and its implementations in finance, but has received little attention in the behavioral economics literature to date.

Finally, we apply our parameter estimates in a model of dynamic trading and find that these improve the model's ability to capture aspects of trading behavior like the disposition effect. In light of our results, we conclude that it is important that future studies in finance incorporate heterogeneity and history dependence in risk preferences, and consider a wider range of preference parameter values than they currently do. For example, further applying our findings to models that have used prospect theory to explain asset pricing anomalies, such as the equity premium and volatility puzzles (e.g., Benartzi and Thaler 1995; Barberis, Huang, and Santos 2001), and examining whether these studies' intuition continues to hold when we relax the assumption of a representative agent with state-independent preferences would be an interesting next step for future research.

Appendix A. Gibbs Sampler

To estimate the model in Section 3, 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 because of the model's complexity we cannot calculate it analytically. Therefore, we use the Markov chain Monte Carlo (MCMC) algorithm to make draws from it. Specifically, 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 resultant 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 for the parameters for each individual and for the parameters' 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. Specifically, using a normal prior with mean $\underline{\kappa}$ and variance $\underline{K}$ for $\mu_{\tilde{\vartheta}}$, its conditional posterior is $\mathcal{N}(\bar{\kappa}, \bar{K})$, that is, 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 3) 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 $\mathcal{W}(\bar{\lambda}, \bar{\Lambda}^{-1})$, that is, 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. We do not have a conjugate prior for the individual-specific parameters, so the conditional posterior is not from a known distribution family, rather we only know its kernel; that is, 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 reject them, balancing their posterior density against their generating density, in an effort to explore the parameter space, while concentrating 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 $\vartheta'_n := g^{-1} \left(g(\vartheta_n^{(k)}) + \sigma L \eta \right)$, 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 reject it and set $\vartheta_n^{(k+1)} = \vartheta_n^{(k)}$.

Appendix 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: we draw (a) the number of bets in the lottery, (b) the type, money staked, and number of individual wagers

for each bet, and (c) the winning odds for each 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, and, 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 and 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, that is, the one with the minimum mean distance to other lotteries in the cluster. The resultant set of lotteries should reasonably “span” the set of randomly drawn lotteries.

References

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

- Abdellaoui, M., H. Bleichrodt, and C. Paraschiv. 2007. Loss aversion under prospect theory: A parameter-free measurement. *Management Science* 53:1659–74.
- Abdellaoui, M., F. Vossman, and M. Weber. 2005. Choice-based elicitation and decomposition of decision weights for gains and losses under uncertainty. *Management Science* 51:1384–99.
- Alvarez, F., L. Guiso, and F. Lippi. 2012. Durable consumption and asset management with transaction and observation costs. *American Economic Review* 102:2272–300.
- Andries, M. 2012. Consumption-based asset pricing with loss aversion. Working Paper, Toulouse School of Economics.
- Andrikogiannopoulou, A., and F. Papakonstantinou. 2016. Heterogeneity in risk preferences: Evidence from a real-world betting market. Working Paper, London School of Economics.
- Barber, B., and T. Odean. 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55:773–806.
- . 2002. Online investors: Do the slow die first? *Review of Financial Studies* 15:455–87.
- Barberis, N. 2012. A model of casino gambling. *Management Science* 58:35–51.
- . 2013. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives* 27:173–96.
- . 2018. Psychology-based models of asset prices and trading volume. In *Handbook of behavioral economics*, eds. D. Bernheim, S. DellaVigna, and D. Laibson, vol. 1, chap. 2, 79–175. Amsterdam, the Netherlands: Elsevier.
- Barberis, N., and M. Huang. 2008. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review* 98:2066–100.
- Barberis, N., M. Huang, and T. Santos. 2001. Prospect theory and asset prices. *Quarterly Journal of Economics* 116:1–53.
- Barberis, N., M. Huang, and R. H. Thaler. 2006. Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review* 96:1069–90.
- Barberis, N., and W. Xiong. 2009. What drives the disposition effect? An analysis of a long-standing preference-based explanation. *Journal of Finance* 64:751–84.
- . 2012. Realization utility. *Journal of Financial Economics* 104:251–71.
- Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum. 2013. The nature of risk preferences: Evidence from insurance choices. *American Economic Review* 103:2499–529.
- . 2018. Estimating risk preferences in the field. *Journal of Economic Literature* 56:501–64.
- Batagelj, V. 1988. Generalized Ward and related clustering problems. In *Classification and related methods of data analysis*, ed. H. H. Bock, 67–74. Amsterdam, the Netherlands: North-Holland.
- Ben-Akiva, M., M. J. 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*, J. Volmuller and R. Hamerslag, eds. Boca Raton, FL: CRC Press.

- Benartzi, S., and R. H. Thaler. 1995. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110:73–92.
- 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* 22:2303–30.
- Booij, A. S., B. M. S. Van Praag, and G. Van De Kuilen. 2010. A parametric analysis of prospect theory's functionals for the general population. *Theory and Decision* 68:115–48.
- 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* 34:315–38.
- Cain, M., D. Law, and D. Peel. 2000. The favourite-longshot bias and market efficiency in uk football betting. *Scottish Journal of Political Economy* 47:25–36.
- Calvet, L. E., J. Y. Campbell, and P. Sodini. 2009. Fight or flight? Portfolio rebalancing by individual investors. *Quarterly Journal of Economics* 124:301–48.
- 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* 107:205–51.
- Chan, Y. L., and L. Kogan. 2002. Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices. *Journal of Political Economy* 110:1255–85.
- Chib, S., and E. Greenberg. 1995. Understanding the Metropolis-Hastings algorithm. *American Statistician* 49:327–35.
- Cohen, A., and L. Einav. 2007. Estimating risk preferences from deductible choice. *American Economic Review* 97:745–88.
- Constantinides, G. M. 1990. Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy* 98:519–43.
- Coval, J. D., and T. Shumway. 2005. Do behavioral biases affect prices? *Journal of Finance* 60:1–34.
- 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*, eds. Y. Ait-Sahalia and L. P. Hansen, vol. 1, 337–82. Amsterdam, the Netherlands: North-Holland.
- Dhar, R., and N. Zhu. 2006. Up close and personal: Investor sophistication and the disposition effect. *Management Science* 52:726–40.
- Dimmock, S. G., and R. Kouwenberg. 2010. Loss aversion and household portfolio choice. *Journal of Empirical Finance* 17:441–59.
- Dorn, D., and P. Sengmueller. 2009. Trading as entertainment? *Management Science* 55:591–603.
- 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* 60:2551–69.
- Feng, L., and M. Seasholes. 2005. Do investor sophistication and trading experience eliminate behavioral biases in finance markets? *Review of Finance* 9:305–51.
- Frazzini, A. 2006. The disposition effect and underreaction to news. *Journal of Finance* 61:2017–46.

- 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* 69:907–46.
- 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* 28:2128–66.
- von Gaudecker, H., A. van Soest, and E. Wengstrom. 2011. Heterogeneity in risky choice behavior in a broad population. *American Economic Review* 101:664–94.
- 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* 6:721–41.
- Gertner, R. 1993. Game shows and economic behavior: Risk-taking on ‘Card Sharks’. *Quarterly Journal of Economics* 108:507–21.
- Glaser, M. 2003. Online broker investors: Demographic information, investment strategy, portfolio positions, and trading activity. Working Paper, University of Mannheim.
- Goeree, M. S. 2008. Limited information and advertising in the U.S. personal computer industry. *Econometrica* 76:1017–74.
- Grinblatt, M., and M. Keloharju. 2001. What makes investors trade? *Journal of Finance* 56:589–616.
- H2 Gambling Capital. 2013. There’s nothing virtual about the opportunity in real-money gambling. Working Paper, H2 Gambling Capital & Odobo, Gibraltar.
- Harrison, G. W., J. A. List, and C. Towe. 2007. Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion. *Econometrica* 75:433–58.
- Hastings, W. 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57:97–109.
- Hey, J. D., and C. Orme. 1994. Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62:1291–326.
- Hoffman, A. O. I. 2007. Individual investors’ needs and the investment professional. *Journal of Investment Consulting* 8:80–91.
- Jullien, B., and B. Salanie. 2000. Estimating preferences under risk: The case of racetrack bettors. *Journal of Political Economy* 108:503–30.
- Kaustia, M., and S. Knupfer. 2008. Do investors overweight personal experience? Evidence from IPO subscriptions. *Journal of Finance* 63:2679–702.
- Kliger, D., and O. Levy. 2009. Theories of choice under risk: Insights from financial markets. *Journal of Economic Behavior & Organization* 71:330–46.
- Kőszegi, B., and M. Rabin. 2006. A model of reference-dependent preferences. *Quarterly Journal of Economics* 121:1133–65.
- . 2007. Reference-dependent risk attitudes. *American Economic Review* 97:1047–73.
- . 2009. Reference-dependent consumption plans. *American Economic Review* 99:909–36.
- Kumar, A. 2009. Who gambles in the stock market? *Journal of Finance* 64:1889–933.

- Kuypers, T. 2000. Information and efficiency: An empirical study of a fixed odds betting market. *Applied Economics* 32:1353–63.
- Kyle, A. S., H. Ou-Yang, and W. Xiong. 2006. Prospect theory and liquidation decisions. *Journal of Economic Theory* 129:273–88.
- Lattimore, P., J. Baker, and A. Witte. 1992. The influence of probability on risky choice: A parametric examination. *Journal of Economic Behavior & Organization* 17:377–400.
- Levitt, S. D. 2004. Why are gambling markets organised so differently from financial markets? *Economic Journal* 114:223–46.
- Li, Y., and L. Yang. 2013. Prospect theory, the disposition effect, and asset prices. *Journal of Financial Economics* 107:715–39.
- 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* 56:606–20.
- Malmendier, U., and S. Nagel. 2011. Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics* 126:373–416.
- Manski, C. F. 2004. Measuring expectations. *Econometrica* 72:1329–76.
- Marschak, J. 1960. Binary-choice constraints and random utility indicators. In *Stanford Symposium on Mathematical Methods in the Social Sciences*, eds. K. J. Arrow, S. Karlin, and P. Suppes, vol. 7, 19–38. Stanford, CA: Stanford University Press.
- McDonald, L. G., and P. Robinson. 2009. *A colossal failure of common sense: The inside story of the collapse of Lehman Brothers*. New York: Crown Business.
- McFadden, D. 1974. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, ed. P. Zarembka, 105–42. New York: Academic Press.
- Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. 1953. Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 21:1087–92.
- Mitton, T., and K. Vorkink. 2007. Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies* 20:1255–88.
- Moskowitz, T. J. 2018. Asset pricing and sports betting. Working Paper, Yale University.
- Odean, T. 1998. Are investors reluctant to realize their losses? *Journal of Finance* 53:1775–98.
- O'Donoghue, T., and C. Sprenger. 2018. Reference-dependent preferences. In *Handbook of behavioral economics*, eds. D. Bernheim, S. DellaVigna, and D. Laibson, vol. 1, chap. 1, 1–77. Amsterdam, the Netherlands: Elsevier Press.
- Pagel, M. 2016. Expectations-based reference-dependent preferences and asset pricing. *Journal of the European Economic Association* 14:468–514.
- Pankoff, L. D. 1968. Market efficiency and football betting. *Journal of Business* 41:203–14.
- 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* 3:137–45.
- . 2009. Sportsbook behavior in the NCAA football betting market: Tests of the traditional and Levitt models of sportsbook behavior. *Journal of Prediction Markets* 3:21–37.
- Polkovnichenko, V. 2005. Household portfolio diversification: A case for rank-dependent preferences. *Review of Financial Studies* 18:1467–502.

- Polkovnichenko, V., and F. Zhao. 2013. Probability weighting functions implied in options prices. *Journal of Financial Economics* 107:580–609.
- Pope, P. F., and D. A. Peel. 1989. Information, prices and efficiency in a fixed-odds betting market. *Economica* 56:323–41.
- 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* 98:38–71.
- Rabin, M. 2000. Risk aversion and expected utility theory: A calibration theorem. *Econometrica* 68:1281–92.
- Routledge, B. R., and S. E. Zin. 2010. Generalized disappointment aversion and asset prices. *Journal of Finance* 65:1303–32.
- Snowberg, E., and J. Wolfers. 2010. Explaining the favorite-longshot bias: Is it risk-love or misperceptions? *Journal of Political Economy* 118:723–46.
- 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* 36:643–60.
- Tversky, A., and D. Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5:297–323.
- Wang, J. 1996. The term structure of interest rates in a pure exchange economy with heterogeneous investors. *Journal of Financial Economics* 41:75–110.
- Weber, M., and C. F. Camerer. 1998. The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization* 33:167–84.
- Woodland, L. M., and B. M. Woodland. 1994. Market efficiency and the favorite-longshot bias: The baseball betting market. *Journal of Finance* 49:269–79.
- Zhang, Y. 2006. Individual skewness and the cross-section of average stock returns. Working Paper, Yale University.

Table 1
Summary statistics of individual and bet attributes

	Mean	Median	SD	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	42	63.93	1	327

This table presents 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.

Table 2
Summary statistics of lottery characteristics

	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
SD	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

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

Table 3
Reduced-form analysis

	A. Participation	B. Lottery choice			
	Duration	Mean	Variance	Skewness	Stake
Constant	0.70 ***	0.37 ***	2.28 ***	0.81 ***	2.47 ***
CumGain	4.61	2.57	22.79	24.85	23.81
	-0.08 ***	0.14 ***	0.12 ***	-0.01	0.14 ***
CumLoss	-6.23	9.05	9.12	-0.87	8.69
	-0.03 ***	0.09 ***	0.09 ***	0.02 **	0.08 ***
CumGain ²	-3.02	6.74	7.30	2.07	6.31
	0.01 ***	-0.00 ***	-0.00 ***	0.00	-0.00 ***
CumLoss ²	4.41	-7.27	-7.67	1.00	-7.23
	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%

Panel A of this table presents results from a panel ordinary least squares (OLS) regression in which the dependent variable is the logarithm of the duration (in days) of each “no play” event—that is, the length of the interval between consecutive play days—for an 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 an 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 an individual over the 7 calendar days preceding the beginning of the “no play” event in panel A (the play day in panel B). All specifications include individual-specific fixed effects. Robust t -statistics are reported below the coefficients. * $p < .1$; ** $p < .05$; *** $p < .01$.

Table 4
Posterior estimates of population mean and variance

A. Means					
	Mean	Median	SD	95% HPDI	NSE
π	0.24	0.24	0.01	[0.22 , 0.27]	0.00
τ	1.36	1.35	0.11	[1.15 , 1.58]	0.01
α	0.86	0.86	0.01	[0.85 , 0.87]	0.00
λ	1.48	1.48	0.04	[1.39 , 1.56]	0.01
γ	0.88	0.88	0.01	[0.87 , 0.90]	0.00
κ	1.41	1.41	0.06	[1.31 , 1.53]	0.01
δ	0.60	0.60	0.02	[0.55 , 0.64]	0.00

B. Variances					
	Mean	Median	SD	95% HPDI	NSE
π	0.05	0.05	0.00	[0.04 , 0.06]	0.00
τ	2.70	2.65	0.39	[2.06 , 3.53]	0.04
α	0.01	0.01	0.00	[0.01 , 0.01]	0.00
λ	0.25	0.24	0.03	[0.20 , 0.32]	0.01
γ	0.01	0.01	0.00	[0.01 , 0.01]	0.00
κ	0.68	0.61	0.18	[0.46 , 1.13]	0.04
δ	0.15	0.15	0.00	[0.14 , 0.16]	0.00

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 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 95% highest posterior density interval (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.

Table 5
Estimated percentiles of model parameters

	Percentiles				
	5th	25th	50th	75th	95th
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

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

Table 6
Comparison of our model with alternatives

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

This table presents the mean and median log-likelihood ($\log L$) across individuals and the deviance information criterion (DIC) for our model and four alternatives: (1) one in which lotteries are chosen randomly; (2) one in which individuals are assumed to be risk neutral, (3) one in which all individuals' preferences are neither heterogeneous nor history dependent, and (4) 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 $\bar{\vartheta} = E[\vartheta|x, y]$, that is, the posterior mean.

Table 7
Disposition effect

μ	<i>A. Without borrowing</i>		<i>B. With borrowing</i>	
	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%

This table presents results on stock market participation and the disposition effect for the two-period portfolio problem presented in Section 4. 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 3), in each panel we present, for several values of the equity risk premium μ (expressed as an annual percentage), the proportions of individuals who 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.

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		

Figure 1
Sample betting slip submitted by an individual
The betting slip shown 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).

Figure 2

The Lattimore, Baker, and Witte (1992) probability weighting function

In this figure, we plot 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.

Figure 3

Value function and probability weighting function

This figure plots the value and probability weighting functions 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 4

Estimated population densities of model parameters

This figure plots 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 . 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 ($\underline{\Delta}=0.1I$) and a high-variance prior ($\underline{\Delta}=10I$), where $\underline{\Delta}$ is the hyperparameter of the prior distribution of Σ_{ϑ} (see Equations (7) and (10) in Section 3.2).

Figure 5
Illustration of the goodness of fit of our model

In this figure, 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 of 100 alternatives has ratio ranking equal to 0.05.

Figure 6

Mean excess realized returns by percentile group of odds

This figure plots the mean excess realized returns for each percentile group of odds in our fixed-odds sports wagering market and in the parimutuel horse race track betting market. For our market, a blue solid line represents the mean excess realized return for each odds group and blue dotted lines represent the 95% bootstrap confidence intervals. For the parimutuel market, we reconstruct the plot in Figure 1 of Snowberg and Wolfers (2010) using 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

This figure illustrates 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 represents our baseline estimated densities. In panel A, the red dash-dotted lines represent 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 represent 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 3.5.

History-Dependent Risk Preferences

Figure 8

Illustration of behavior

This figure illustrates the behavior of three benchmark individuals facing 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 represent the individual's value function; the black dots denoted by L_1 and G_1 represent 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.

Figure 9

Stock market participation and the disposition effect

This figure shows results for stock market participation and the disposition effect for the two-period portfolio problem of Section 4, 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 .