

# MODELING PREFERENCE

## HETEROGENEITY WITHIN AND ACROSS BEHAVIORAL TYPES: EVIDENCE FROM A REAL-WORLD BETTING MARKET\*

Angie Andrikogiannopoulou  
King's College London

Filippos Papakonstantinou  
King's College London

### Abstract

While it is commonly accepted that risk preferences differ across individuals, studies that estimate them typically allow for limited heterogeneity. We develop a methodology that allows for richer representation of heterogeneity both within and across utility types characterized by different behavioral features. This enables us to improve individual- and population-level estimates, and to assess the relative importance of loss aversion and probability weighting, and their prevalence in the population. Applying our model to individual sports-betting choices, we find that utility curvature alone does not explain observed choices and, while two-thirds of individuals exhibit loss aversion, all exhibit probability weighting.

*Keywords: Prospect Theory, Loss Aversion, Probability Weighting,  
Hierarchical Mixture, Bayesian Methods*

*JEL Classification: D90, C11, C51, D81, D12, L83*

---

\*For helpful comments, the authors thank Adrian Bruhin, Thomas Epper, Bo Honore, Andriy Norets, Chris Rogers, Bernard Salanie, Paolo Sodini, Hyun Shin, Peter Wakker, Wei Xiong, as well as participants at the 2014 European Economic Association Annual Congress and the 2018 Center for Economic Analysis of Research Workshop, and seminar participants at Cambridge University, Einaudi Institute for Economics and Finance, Imperial College London Business School, McGill University, Oxford University, Princeton University, the Stockholm School of Economics, the University of Amsterdam, the University of Geneva, and the Wharton School. This work was supported by an Amazon Web Services in Education research grant award. This work has previously been circulated under the title "Heterogeneity in risk preferences: Evidence from a real-world betting market". 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.

Since Kahneman and Tversky (1979) proposed prospect theory, a plethora of experimental and field studies have shown that individual behavior deviates from the predictions of expected utility theory and exhibits patterns consistent with loss aversion and probability weighting. But even though prospect theory has become the most widely accepted behavioral theory of choice, modeling limitations and the limited availability of rich data, particularly in the field, have made it difficult to distinguish between the different prospect theory features and assess their relative importance and prevalence in the population. Hence, an important question remains unresolved: which of prospect theory’s features are most useful in explaining behavior and how prevalent are they in the population? In this paper, we aim to address this question in the field by (1) proposing a hierarchical mixture model that allows for an extensive representation of individual heterogeneity *within* and *across* utility types characterized by the different prospect theory features, (2) designing a Bayesian algorithm to estimate this model, and (3) utilizing a panel dataset of individual sports betting choices that is well-suited to distinguishing the prospect theory features.

Methodologically, we build on the idea that preferences are drawn from a population distribution hence that we can learn about the preferences of an individual by taking into account the choices of *all* individuals. Consistent with this idea, some extant studies use hierarchical models in which prospect theory parameters are drawn from a single normal distribution. But this would be inappropriate if individuals cluster into distinct sub-populations that exhibit some behavioral features but not others. Other studies consider a mixture of utility types with different behavioral features but with no variation across individuals within types. However, this can result in misleading inferences if individual behavior departs from that of the group average. We generalize these approaches by introducing preference heterogeneity along two dimensions. First, we model the population as a mixture of four sub-populations that represent four prospect-theory types: a type that exhibits neither probability weighting nor loss aversion, two types that each turn on one of these features, and an unrestricted type that exhibits both features. Second, we incorporate preference heterogeneity within each type by allowing the preference parameters of individuals belonging in each type to be drawn from a sub-population distribution. This hybrid model produces statistical inference about the prevalence of the behavioral features in the population,<sup>1</sup> while forgoing the restrictive assumption that there is a representative agent for each type.

To estimate our model, we design a Bayesian Markov chain Monte Carlo algorithm which enables us to jointly estimate the proportions of each behavioral type, the entire mixture distribution of the population-level preference parameters, as well as the individual-level preference parameters. Thus, our methodology shares information across individuals in a flexible manner dictated by the data and so improves the accuracy of individual- and population-level estimation, as what we learn about each individual’s preferences and type allocations feeds into the shape of the population distribution, and vice versa. In a simulation analysis, we illustrate that ignoring heterogeneity within or across types can severely misestimate the distribution of preference parameters, misallocate as much as half the population from fully behavioral to fully rational, and thus produce misleading implications about individual and aggregate behavior. Notably, our methodology can be applied to various settings in economics in which substantial heterogeneity is likely to exist both within and across groups.

---

<sup>1</sup>This can be viewed as an analogue of multiple testing whereby multiple hypothesis tests (one for each individual) of the nulls of no loss aversion and no probability weighting are performed.

We use our methodology to study individual choices in a sports betting market. Sports betting is a very common and economically important activity, with “half of America’s population and over two-thirds of Britain’s [placing a] bet on something” per year and with \$1 trillion wagered on sports, globally, per year.<sup>2</sup> Furthermore, betting markets provide an advantageous natural laboratory for risk preference estimation. First, they allow for a lottery representation of choices because prizes (i) are determined exogenously by match outcomes and (ii) are associated with probabilities that are predicted accurately by the quoted odds; there is no adverse selection or moral hazard to complicate matters. Second, they allow the estimation of all prospect theory features, as choices consist of a wide range of prizes and probabilities, including gains and losses as well as small and large stakes. Third, they allow the estimation of preference heterogeneity with improved accuracy, as individuals in this setting typically make a relatively large number of choices.

Our model of behavior is based on the cumulative prospect theory (CPT) of Tversky and Kahneman (1992), according to which individuals use a value function defined over gains and losses relative to a reference point, have different sensitivity to losses and gains, and systematically distort event probabilities by weighting the gain/loss cumulative distribution. We find that utility curvature alone does not adequately explain choices, and the distinctive features of prospect theory—loss aversion and probability weighting—offer a significant improvement. But, importantly, our study reveals that these two features are not equally prevalent. Individuals can be partitioned into two groups: two-thirds exhibit both behavioral features, and one-third exhibit probability weighting but *not* loss aversion. This indicates that, contrary to commonly accepted wisdom, prospect theory should not be viewed as a monolithic theory with a set of features that all individuals exhibit, but rather reality is more nuanced. This more nuanced view could facilitate and guide the use of prospect theory in applied models. For example, our finding that probability weighting plays a more important role than loss aversion could be useful in finance, where loss aversion has sometimes led to counterfactual predictions for individual investor behavior and for the aggregate stock market (see Barberis, 2013 for a review).

Even though the group of people who bet on sports is quite large, as with other field studies on risk preference estimation our results are not immune to a potential sample selection bias. For example, very loss-averse individuals may be less likely to engage in sports betting, hence our estimates may not be representative of the whole population. Nonetheless, it is reassuring that (i) our estimated loss aversion is only weakly related to the propensity to accept betting opportunities, and (ii) all our preference parameter estimates are not only very heterogeneous but also within the range of those reported in the experimental literature (see Wakker, 2010). Furthermore, our estimates make realistic predictions when used in a different economic situation. Applying our estimated preference parameters in a standard portfolio choice problem, we find that they yield optimal portfolios similar to those observed in surveys on household portfolio allocations. All this indicates that a potential sample selection bias is likely mild. Regardless, keeping in mind that extrapolating preference parameters from *any* group or context to another necessitates additional assumptions, our estimates from this setting are complementary to those from other settings (e.g., insurance choice), and may even be particularly relevant for sub-populations or contexts that share similar characteristics with those in wagering markets (e.g., traditional financial markets).

---

<sup>2</sup>See *The Economist*, 8 July 2010 and the H2 Gambling Capital report of 2013.

Following our main analysis, we explore extensions and robustness checks. Regarding individuals' beliefs, our baseline is the standard assumption in the literature that they are rational which, in our setting, can be approximated quite well by the probabilities implied by the market prices. In an extension, we explore two alternative beliefs specifications: one that uses a different approximation of rational beliefs and one that allows for heterogeneous subjective beliefs that deviate from the rational. We are able to identify risk preferences while allowing for irrational beliefs by exploiting the great variety in the choices available in a sportsbook. For example, even if someone believes an outcome is mispriced (e.g., that Nadal's edge over Federer is undervalued), he still needs to choose an event (e.g., to back Nadal to win, or to win by a lot) as well as whether to combine this wager with others. These choices correspond to different risk profiles, so are informative about risk preferences even if subjective probabilities deviate from the rational. Our results remain similar to those from the main analysis, with considerable heterogeneity in preferences both within and across utility types. We also show that observed bets are unlikely to be driven by (real or perceived) superior information or affinity toward specific teams (e.g., home-area teams), since almost all individuals place wagers on a large variety of leagues and teams rather than any specific one, and virtually no individuals generate significantly positive returns from betting. Regarding the choice sets individuals face, our baseline assumption is that individuals consider all lotteries available in the sportsbook, except those that have rarely been chosen by anyone in our sample. In a robustness check, we consider an alternative choice set containing only lotteries that respect a tight, but reasonable, budget constraint of a few hundred euros.

Our work is most directly related to the small number of studies that compare utility types and/or stress the importance of probability weighting. Using experimental data, Bruhin, Fehr-Duda and Epper (2010) consider a mixture of utility types, each with different parameters but with no variation across individuals within types, and find that 80% of the individuals exhibit some probability weighting. Conte, Hey and Moffatt (2011) also use a mixture of utility types, but they first estimate the types' parameters assuming *all* individuals belong to *both* types and they subsequently allocate individuals to these types. We estimate population- and individual-level parameters simultaneously and in a consistent manner, as what we learn about individuals' preferences and type allocations feeds into the estimation of the corresponding population parameters only, and vice versa. In the field, Snowberg and Wolfers (2010) compare two models—one with risk aversion but no probability weighting and one with probability weighting but risk neutrality—and find that the latter better explains aggregate prices in the horse-racetrack betting market. Rather than allocate an entire market to a single restrictive type, we use individual-level data and we allocate different individuals to different types. Barseghyan et al. (2013) find that some kind of “generic” probability distortion is important in explaining insurance deductible choices; they do not juxtapose probability weighting with loss aversion because, with their data, the two cannot be distinguished.

More generally, our paper contributes to the fundamental literature on preference estimation. The majority of evidence on estimating preferences with prospect-theory features is from lab experiments (e.g., Kahneman and Tversky, 1979, von Gaudecker, van Soest and Wengstrom, 2011) and TV game shows (e.g., Post et al., 2008). Some studies estimate preferences in the field, mostly using aggregate prices to infer the preferences of a representative agent (e.g., Jullien and Salanie, 2000, Polkovnichenko and Zhao, 2013), while only a handful estimate heterogeneous preferences using individual-level data: Barseghyan et al.

(2013) study deductible choices in insurance markets and Andrikogiannopoulou and Papakonstantinou (2020) study sport betting choices, as we do, but focus on history dependence and only allow for limited heterogeneity across individuals, which produces qualitatively distinct results; our paper is different in that it accommodates rich heterogeneity both within and across utility types. Overall, the combination of our hierarchical mixture model and our data enables us to improve estimation accuracy and to assess—for the first time—the relative importance and prevalence of the behavioral features of prospect theory.

Our study is also related to the small literature that focuses on heterogeneity in beliefs and questions the importance of heterogeneity in risk preferences: Gandhi and Serrano-Padial (2014) assume risk neutrality and explain aggregate betting prices using heterogeneity around the rational beliefs, and Ericson et al. (2021) assume CARA utility and show that heterogeneity in (perceived) risk exposure is necessary and sufficient to explain the health insurance choices they observe. While in our baseline analysis we follow the main literature in assuming rational beliefs (see Barseghyan et al., 2018 for a review), in one of our extensions we allow subjective beliefs to deviate from the rational and to vary across individuals. In this analysis we show that, even with heterogeneous beliefs, heterogeneity in risk preferences remains important.

In Section 1, we present our model of individual behavior, our mixture model of heterogeneity, and our estimation methodology. In Section 2 we present the data and in Section 3 our results. In Section 4 we use our estimates in a portfolio choice application, and in Section 5 we present robustness checks. In Section 6, we conclude. An Appendix contains additional details on the estimation procedure, and an Internet Appendix contains additional tables and figures supporting our analysis.

## 1 Model and estimation

### 1.1 Model of individual behavior

Consider an individual who at time  $t$  has the opportunity, with probability  $\pi$ , to choose among a set of risky prospects. In our sports betting context,  $\pi$  could capture how busy he is, hence how frequently he logs on the sportsbook. Given this opportunity, the individual either rejects it and chooses the safe lottery  $\delta_0$  that pays 0, so we do not observe a choice, or he accepts it and chooses a lottery that we observe.

We adopt the random-utility model and decompose the utility that the individual with preference parameters  $\rho$  obtains from choosing lottery  $L_j$  on day  $t$  into

$$U(\rho; L_j) + \varepsilon_{jt},$$

with  $U$  a deterministic utility and  $\varepsilon_{jt}$  a stochastic component that accounts for unmodeled factors and is necessary to reconcile observed behavior with a(ny) deterministic choice theory. As is common in the literature,  $\varepsilon$  is assumed i.i.d. drawn from the double exponential with location normalized to 0 and inverse scale  $\tau > 0$ .<sup>3</sup>

---

<sup>3</sup>While the random-utility model is very commonly used in the literature in risk preference estimation, recently Apesteguia and Ballester (2018) have argued that its use may be problematic as it may pose identification problems, and instead they suggest the use of a random-parameters model. However, both the partial identification approach to classical inference as well as the Bayesian approach to inference (which we employ here) recognize that identifiability need not have a binary answer, but rather model parameters are identified to the extent that the data contains information about them. As such, the choice between the random-utility vs. the random-parameter model should not be dictated by concerns relating to estimation but rather by their relative merits in describing observed behavior. We do not take a stand on this debate, rather we simply follow the majority of the literature and use the random-utility model, taking care to check that our data does indeed contain information about the model parameters.

Letting  $\mathbf{y} := \{y_t\}$  be the sequence of lottery choices we observe and  $\mathbf{x} := \{x_t\}$  the sequence of dummies indicating whether we observe a choice at time  $t$ , the likelihood of observing  $(\mathbf{x}, \mathbf{y})$  is

$$p(\mathbf{x}, \mathbf{y} | \pi, \tau, \rho) = \prod_t \left\{ x_t \pi p(y_t | \tau, \rho) + (1 - x_t) \left[ \pi p(y_t = \delta_0 | \tau, \rho) + (1 - \pi) \right] \right\}, \quad (1)$$

where  $p(y_t | \tau, \rho)$  is the probability that individual with parameters  $\tau, \rho$  chooses lottery  $y_t$ . An individual chooses  $y_t = L_j$  from choice set  $C$  with probability

$$\begin{aligned} p(y_t = L_j | \tau, \rho) &= p\left(U(\rho; L_j) + \varepsilon_{jt} > \max_{i \neq j \in C} \{U(\rho; L_i) + \varepsilon_{it}\}\right) \\ &= \frac{\exp(\tau U(\rho; L_j))}{\sum_{i \in C} \exp(\tau U(\rho; L_i))}, \end{aligned} \quad (2)$$

which follows from the distribution of  $\varepsilon$ .

Utility  $U$  represents Kahneman and Tversky's (1992) cumulative prospect theory (CPT), which has been shown to do well in explaining behavior in various settings, including gambling (e.g., Barberis, 2012). Specifically, an individual with preference parameters  $\rho := (\alpha, \lambda, \gamma, \kappa)'$  evaluates lottery  $L$  with prizes  $z_{-m} \leq \dots \leq z_0 = 0 \leq \dots \leq z_M$  and corresponding probabilities  $\{p_i\}$  by comparing final to current wealth and assigning to the lottery the value

$$U(\rho; L) := \sum_i w_i u(z_i),$$

where  $u$  is the value function and  $\{w_i\}$  are the decision weights. The value function is

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

Parameter  $\alpha > 0$  measures value function curvature, where  $\alpha < 1$  implies risk aversion (seekingness) over gains (losses) while  $\alpha > 1$  implies risk seekingness (aversion) over gains (losses). The relative sensitivity to gains and losses is measured by  $\lambda > 0$ , where  $\lambda > 1$  implies loss aversion and  $\lambda < 1$  implies gain seekingness. Though many early studies have found that, on average,  $\alpha \leq 1$  and  $\lambda \geq 1$ , we do not impose these commonly used restrictions, because there is evidence that some individuals exhibit risk and/or gain seekingness (e.g., Abdellaoui, Bleichrodt and Paraschiv, 2007; von Gaudecker, van Soest and Wengstrom, 2011).<sup>4</sup>

The decision weight  $w_i$  takes the form

$$w_i = \begin{cases} w(p_i + \dots + p_M) - w(p_{i+1} + \dots + p_M) & \text{for } 0 \leq i \leq M \\ w(p_{-m} + \dots + p_i) - w(p_{-m} + \dots + p_{i-1}) & \text{for } -m \leq i < 0 \end{cases}, \quad (4)$$

where the probability weighting function  $w$  has the flexible Goldstein and Einhorn (1987) form:

$$w(p; \gamma, \kappa) := \frac{\kappa p^\gamma}{\kappa p^\gamma + (1 - p)^\gamma}. \quad (5)$$

The curvature of  $w$  is measured by  $\gamma > 0$ , where  $\gamma < 1$  ( $\gamma > 1$ ) corresponds to an inverse-S shape (S shape), so overweighting (underweighting) of extreme outcomes. Elevation is measured by  $\kappa > 0$ ; e.g., with  $\gamma = 1$ ,

<sup>4</sup>In our estimation, we find that few individuals have  $\alpha > 1$  and/or  $\lambda < 1$ . Imposing the restrictions  $\alpha \leq 1$  and  $\lambda \geq 1$  yields very similar results to the ones from the unrestricted estimation.

$\kappa > 1$  corresponds to a concave  $w$  so overweighting of outcomes in the right (left) tail of the distribution of positive (negative) outcomes, and  $\kappa < 1$  corresponds to a convex  $w$ .

To reduce model complexity and facilitate identification, we assume the following. First,  $\alpha$  and  $\gamma$ , hence the value function and the probability weighting function, are the same for gains and losses. Indeed, empirical studies usually estimate similar values for these parameters over gains and losses (e.g., Tversky and Kahneman, 1992; Abdellaoui, 2000), so it is common in the literature to make this assumption (e.g., Post et al., 2008; Baillon, Bleichrodt and Spinu, 2020). Furthermore, to assume different value function curvatures for gains and losses is problematic both theoretically, as it implies loss aversion is not well defined (Kobberling and Wakker, 2005; Wakker, 2010), and empirically, as it hinders the identification of loss aversion (Nilsson, Rieskamp and Wagenmakers, 2011). Second, the reference point individuals use to separate gains from losses equals current wealth. While Kahneman and Tversky (1979) suggest that the status quo is a natural choice for most choice situations, other possible reference points such as an expectation have been proposed in the literature. But, as the empirical evidence on the characterization of the reference point is mixed (see O’Donoghue and Sprenger, 2018 for a review), the status quo remains the most common assumption in the literature. Indeed, in a recent experiment, Baillon, Bleichrodt and Spinu (2020) find that the status quo is one of the most commonly used reference points.

## 1.2 Mixture model of preference heterogeneity

We introduce heterogeneity in preferences by proposing that parameters  $\rho_n$  for individual  $n$  are drawn from a population distribution. Since the elements of  $\rho_n$  are bounded, we model heterogeneity in terms of transformations  $\tilde{\rho}_n := g(\rho_n)$ , where  $g$  represents Johnson transformations that map elements of  $\rho_n$  to  $(-\infty, +\infty)$ ; e.g., for  $\theta \in [\underline{\theta}, \bar{\theta}]$ , we define  $\tilde{\theta} := \ln((\theta - \underline{\theta})/(\bar{\theta} - \theta))$ .

Previous studies that estimate heterogeneous preferences with information sharing model heterogeneity by assuming the (transformed) parameters are drawn from a normal with population mean  $\mu_{\tilde{\rho}}$  and variance  $V_{\tilde{\rho}}$ . That is, they propose  $p(\tilde{\rho}_n | \mu_{\tilde{\rho}}, V_{\tilde{\rho}}) = f_{\mathcal{N}}(\tilde{\rho}_n | \mu_{\tilde{\rho}}, V_{\tilde{\rho}})$ , or equivalently

$$p(\rho_n | \mu_{\tilde{\rho}}, V_{\tilde{\rho}}) = f_{\rho}(\rho_n | \mu_{\tilde{\rho}}, V_{\tilde{\rho}}), \quad (6)$$

where  $p(\cdot | \cdot)$  denotes a conditional density,  $f_{\mathcal{N}}$  is the normal density, and  $f_{\rho}$  follows from the Jacobian. This specification allows for individual heterogeneity within a utility type, but is restrictive as it requires that *all* individuals exhibit *all* features of behavior associated with this utility. Here we go further and allow for heterogeneity not only *within* but also *across* utility types. This generalization (i) accommodates a flexible representation of preference heterogeneity hence shares information across individuals more objectively than a restrictive model, and (ii) enables us to produce statistical inference about the proportions of individuals exhibiting different behavioral features, and thus to evaluate the relative importance of each feature in the population.

To this end, we relax the assumption that preferences are drawn from a single distribution, and instead propose that they are drawn from a mixture. Thus, Equation 6 becomes

$$p(\rho_n | \{\mu_{\tilde{\rho}}^q, V_{\tilde{\rho}}^q\}) = \sum_{q \in Q} h^q f_{\rho}^q(\rho_n | \{\mu_{\tilde{\rho}}^q, V_{\tilde{\rho}}^q\}), \quad (7)$$

where  $Q$  is the set of utility types;  $h^q \in [0, 1]$  is the proportion of type  $q$  in the population; and  $f_\rho^q(\cdot | \mu_\rho^q, V_\rho^q)$  is the density of the parameters for type  $q$ . That is, each individual belongs to type  $q$  with probability  $h^q$  and, conditional on belonging to it, his parameters are drawn from  $f_\rho^q$ ; we retain the convenient assumption that transformed parameters are normally distributed within a type, unless they are restricted (e.g.,  $\lambda = 1$  for individuals in the type with no loss aversion). It is also useful to define unobserved indicators  $e_n^q \in \{0, 1\}$  that indicate an individual's type. Letting  $\mathbf{h} := \{h^q\}$ , then  $\mathbf{e}_n := \{e_n^q\}$  has the multinomial  $\mathcal{M}(1, \mathbf{h})$  distribution.

To assess the importance of loss aversion and probability weighting in explaining behavior, we estimate a mixture of four utility types: (i) a type that exhibits neither feature, (ii) a type that only turns off loss aversion, (iii) a type that only turns off probability weighting, and (iv) an unrestricted type that exhibits all CPT features. If individual  $n$  belongs to the type with neither loss aversion nor probability weighting,  $\rho_n := (\alpha_n, \mathbf{1})'$ . If he has no loss aversion,  $\rho_n := (\alpha_n, 1, \gamma_n, \kappa_n)'$ . If he has no probability weighting,  $\rho_n := (\alpha_n, \lambda_n, \mathbf{1})'$ . If he belongs to the unrestricted type,  $\rho_n$  is unconstrained.

We note that, as with the preferences, the inverse scale  $\tau$  of the stochastic utility component and the probability  $\pi$  of having the opportunity to make a risky choice are individual-specific.

### 1.3 Bayesian estimation

We use Bayesian estimation to combine information in the data—individuals' lottery choices and choice frequencies—with priors and produce posterior inference. Since we have no prior information about the model parameters, we use weak priors. Specifically, we utilize a two-layer hierarchical prior that models the lack of prior information through a distribution over priors. The first layer is essentially the heterogeneity model presented above, i.e., individual parameters are drawn from a population distribution with unknown parameters estimated from the data. The second layer is a hyperprior for these population parameters. This approach contains data-driven information, so allows for a more objective and robust inference (Robert, 2007). Specifically, for the population proportions  $\mathbf{h}$ , we use a Dirichlet prior  $\mathbf{h} \sim D(\underline{\mathbf{h}})$  with  $\underline{\mathbf{h}} = \mathbf{1}$ , which implies an uninformative uniform distribution over all possible values of  $\mathbf{h}$ . For the population means and variances, we use the independent Normal-inverse-Wishart prior, i.e.,  $(\mu_{\tilde{\pi}\tilde{\tau}}, V_{\tilde{\pi}\tilde{\tau}}) \sim \text{NIW}(\underline{\boldsymbol{\kappa}}, \underline{\mathbf{K}}, \underline{\zeta}_{\tilde{\pi}\tilde{\tau}}, \underline{\mathbf{Z}})$  and  $(\mu_\rho^q, V_\rho^q) \sim \text{NIW}(\underline{\boldsymbol{\kappa}}, \underline{\mathbf{K}}, \underline{\zeta}_\rho^q, \underline{\mathbf{Z}})$ , with parameters that render the priors weak. In Section 5.4, we show that our estimates are robust to alternative priors.

The joint posterior density is proportional to the likelihood times the joint prior, but we cannot calculate it analytically. Thus, we obtain information about it using a Gibbs sampler, which partitions the parameters into blocks such that it is possible to draw sequentially from the posterior of each block conditional on the other blocks and the data. These sequential draws form a Markov chain that, under weak conditions, converges to the joint posterior regardless of the initial draw (see Geweke, 1999).<sup>5</sup> For example, in our model, individuals' allocations  $e_n$  to types are a sample from the multinomial  $\mathcal{M}(1, \mathbf{h})$  distribution, so they can be used to learn about the population proportions  $\mathbf{h}$ . Furthermore, for individuals belonging to type  $q$ , their parameters  $\tilde{\rho}_n$  are a sample from  $\mathcal{N}(\mu_\rho^q, V_\rho^q)$ , so we can use them to learn about  $\mu_\rho^q$  and  $V_\rho^q$ .

<sup>5</sup>Intuitively, this estimation solves for an equilibrium density—the joint of the model parameters—which is roughly a stochastic analogue of a fixed point.

A non-trivial complication arises because the four type-specific distributions from which individuals' preference parameters  $\rho_n$  are drawn overlap at sets of measure zero. As a result, we cannot draw sequentially from the conditional posteriors of the type allocations  $e_n$  and of the  $\rho_n$ , because it would be impossible to explore the entire posterior of  $\rho_n$ . To circumvent this, we design a Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) that enables us to draw in one step from the *joint* conditional posterior of  $e_n, \rho_n$ . In this step, we draw candidates by first drawing the allocation to a utility type and then, conditional on the allocation, drawing the preference parameters. These candidates are accepted or rejected so that the Markov chain moves toward parameters with high posterior probability. Intuitively, we combine information in each individual's likelihood with population-level information—the proportions and the distribution of the preference parameters for each type—to learn (i) about the probability that the individual belongs to each type and (ii) about the distribution of his preference parameters conditional on belonging to each type. For example, consider how an individual is allocated to the  $\lambda = 1$  type with no loss aversion. If (i) this type's proportion in the population is high, (ii) the individual's choices are better explained by  $\lambda_n = 1$  hence his likelihood at  $\lambda = 1$  is high, and (iii) the overlap between this likelihood and the population distribution at  $\lambda \neq 1$  is low, then the individual likely belongs to the  $\lambda = 1$  type. Conversely, the higher is the probability that an individual belongs to the  $\lambda = 1$  type, the higher is the posterior estimate for this type's proportion in the population.

## 2 Data and lottery construction

### 2.1 Sports betting data

We apply our model and estimation to a dataset of individual activity in a sports betting market. A sportsbook offers betting opportunities on a variety of sports (e.g., soccer, tennis), fixtures (e.g., English soccer league matches), and events (e.g., the match winner or the exact score). The bookmaker sets the odds for each possible outcome of an event, i.e., the inverse of the price for receiving a unit payoff in case the outcome is realized. Customers place wagers at these odds and the bookmaker takes the opposite side with the intention of making a profit through a commission. For example, an individual betting on the match outcome can select one of three outcomes—home win, draw, or away win. If the outcomes' prices are 0.55, 0.30, and 0.20 (quoted as having odds  $1/0.55 \approx 1.82$ ,  $1/0.30 \approx 3.33$ , and  $1/0.20 = 5.00$ ) then a wager backing the home team to win will net €82 for each €100 staked if it wins. The sum of the prices of all event outcomes is the price of a portfolio that pays €1 for sure, and the amount by which this sum exceeds 1 is the bookmaker's commission per euro wagered. The quoted prices are associated with the outcomes' *implied probabilities*, calculated in this example as  $0.55/(0.55+0.30+0.20) \approx 0.52$ ,  $0.30/(0.55+0.30+0.20) \approx 0.29$ , and  $0.20/(0.55+0.30+0.20) \approx 0.19$ , respectively. An individual can place a single wager or combine multiple wagers into various bet types. Specifically,  $N$  wagers can be combined in a type- $N$  *accumulator* bet, which yields gross return equal to the product of the  $N$  wagers' gross returns if all of them win, and zero otherwise. Accumulator bets can further be combined into a *permutation* bet that includes the  $\binom{N}{k} = \frac{N!}{k!(N-k)!}$  type- $k$  accumulators for some  $k \leq N$ , or a *full-cover* bet that includes the  $\sum_{l \leq k \leq N} \binom{N}{k}$  accumulators of type  $l$  through  $N$ , for  $l \in \{1, 2\}$ . Finally, the individual chooses how much money to stake, divided evenly among all accumulators.

Our data consists of the betting histories of 336 randomly chosen customers of a European online sportsbook over the period 2005–2010, also used in Andrikogiannopoulou and Papakonstantinou (2020). For each bet placed, we observe: date, event bet and outcome chosen (e.g., Chelsea to defeat Arsenal), amount wagered, bet type, quoted odds, and bet result. In Table 1, we provide summary statistics for the characteristics of the individuals we observe and their selected bets. The mean age is 33 years and most individuals (93%) are male. The average individual places wagers on 35 days, at a frequency of once a week, and spreads his bets across 35 leagues, 183 teams, and 23 event types (final outcome, exact score, etc.). Furthermore, in Panel a (Panel b) of Figure 1 we plot a histogram of the percentage of bets by each individual on the league (team) that he has bet most frequently. We see that most individuals place a small proportion—less than 20% (3%)—of their bets on any specific league (team). This suggests it is unlikely that choices are driven by a systematic preference or information about any specific league or team.

Finally, in Figure 2 we examine if the probabilities implied by the quoted odds are close to the true probabilities. We collect from the Web the closing odds and final outcomes of all soccer matches available in our sportsbook.<sup>6</sup> Then we divide outcomes into 100 odds percentiles, and for each group we plot the mean probability implied by the odds against the actual win frequency. Even though the implied probabilities of favorite outcomes are slightly below their win frequencies, indicating a small favorite-longshot inefficiency, we generally observe that deviations of the implied from the objective win probabilities in this market are small. So, as in other fixed-odds sports-betting markets (Pope and Peel, 1989; Woodland and Woodland, 1994), the bookmaker under study finds it optimal to set odds close to the efficient ones rather than the ones that “balance” the money wagered on each outcome. Consistent with this, in Section 5.3 we find that no individual in our data realizes excess returns statistically different from zero.

## 2.2 Bets as lottery choices

To estimate preferences using field data, it is necessary to map the complexity of the real world to a manageable framework. In our betting context, we adopt assumptions that are similar in spirit to the assumptions used by the other studies that estimate preferences in the field (see, e.g., studies on insurance and game-show choices by Cohen and Einav, 2007 and Post et al., 2008, respectively).

First, we create a lottery representation of the simplest bet, which is to select an outcome in one event. This bet has two potential prizes: if the selected outcome occurs, the prize is the stake times the selection’s net return, else the stake is lost. But, as in any field setting, individuals’ subjective beliefs with respect to the probabilities of these prizes are unobservable, giving rise to an identification problem which is commonly resolved by assuming beliefs are rational; see Manski (2004) for a general discussion of this issue, and Barseghyan et al. (2018) for a discussion specifically in the context of preference estimation.<sup>7</sup>

---

<sup>6</sup>In this analysis we only use soccer matches because data on their outcomes are more readily available than for other sports. Furthermore, soccer betting constitutes the most popular market segment in the sportsbook we study.

<sup>7</sup>For example, studies using horse-racetrack betting data approximate each outcome’s true probability with the win frequency of outcomes with similar prices; studies using insurance choices assume that each individual holds rational beliefs regarding his probability of incurring a claim during the policy period, and estimate these beliefs by regressing claim rates on demographics; and game show studies assume contestants’ beliefs are consistent with a rule that explains the majority of the variance of the presenter’s strategic behavior. A couple of exceptions that assume non-rational beliefs (but restrict preferences) are Gandhi and Serrano-Padial (2014) and Ericson et al. (2021).

In our main analysis, we follow the majority of the literature and assume beliefs are rational which, in our setting, can be approximated quite well by the probabilities implied by the market prices.<sup>8</sup> Importantly, two features of our data indicate that subjective probabilities are likely to be close to the implied. First, several online tools (which can be accessed through the betting platform) help individuals calculate the implied probabilities from the quoted prices, so they are observable by the individuals themselves when they place their bets. As a result, it is likely that the implied probabilities serve as a focal point that anchors individuals' beliefs as well as discourages significant deviations from them. Second, as discussed above, individuals place wagers on a large variety of leagues and teams. So, given how infrequently individuals place bets on specific leagues or teams, it is unlikely that they have superior information about any of them. Nevertheless, in a sensitivity analysis in Section 5.1, we exploit the great variety in the available choices in our setting to identify risk preferences while allowing for small but significant deviations of the subjective beliefs from the rational beliefs and we show that our results remain very similar.

Subsequently, we create lottery representations of accumulator bets and of more complex bets by combining the lotteries of the constituent accumulators.<sup>9</sup> Once we have a lottery representation for each bet, we create a lottery representation for each play session—defined as a single day—for each individual;<sup>10</sup> we represent this day lottery by constructing all possible combinations of payoffs from all the bets placed on a given day. We drop 1,043 day lotteries which contain bets on related events hence cannot easily be constructed, and our final sample contains 11,490 lotteries.<sup>11</sup> In Table 2, we report summary statistics for the characteristics of these chosen day lotteries, pooled across all individuals. About half of the lotteries contain 2 prizes, but 25% of lotteries contain at least 6 prizes. The bet amount ranges from 1 cent to thousands of euros (€5,500), while the maximum prize ranges from 1 cent to hundreds of thousands of euros. Wagers have negative mean return, because of the commission incorporated in the prices, and most lotteries have positive skewness. Overall, the summary statistics demonstrate that there is substantial variation in the lotteries we observe.

Next, we specify the set of lotteries from which individuals make their choices. Similar to other field studies (e.g., Barseghyan et al., 2013), we assume that the individuals' consideration set contains all available choices except those rarely chosen by anyone. To construct this set, we generate many lotteries by drawing their risk characteristics from their empirical distributions. The procedure we use resembles the one individuals use to place wagers: we select the number of bets per session, the type of each bet, the odds associated with each wager, and the amount staked. Then, we augment this set with the chosen lottery and the safe lottery which corresponds to not placing a bet. This choice set generation is a computationally convenient reduced-form analogue of a structural model in which the more complex and costly a lottery is,

---

<sup>8</sup>In an extension, we show that our results do not change if we approximate rational beliefs using the win frequencies of outcomes with similar prices, which—as discussed above—are close to the implied probabilities.

<sup>9</sup>In the Internet Appendix, we present examples that illustrate how we create the lottery representations of individuals' bets.

<sup>10</sup>Alternatively, a play session could be defined, at one extreme, as an individual bet, and at the other extreme as the portfolio of all outstanding bets. The former is too narrow since individuals often place multiple bets at the same time, while the latter is in fact very close to how we have defined the play session as a single day, as most bets in our sample have been placed on the same day as the event.

<sup>11</sup>When multiple bets are placed on related events (e.g., the winner of a match and its exact score), we cannot construct the lottery without further information about the outcomes' joint distributions, so we drop days that involve related bets. These days are relatively few, and it is unlikely that choices on these days are systematically different than on others, so dropping them should not affect our results.

the less likely it is to be in the choice set; for similar reduced-form approaches, see Ben-Akiva et al. (1984) and Brownstone, Bunch and Train (2000). In Section A of the Appendix, we present more details on the choice set generation procedure, and in Section 5.2 we consider an alternative choice set that accounts for the possibility that individuals may face a reasonable budget constraint of a few hundred euros.

As mentioned above, each individual places wagers on a great variety of leagues, teams, and events, indicating that his choices are not driven by a systematic preference for specific teams or events. As a result, even though different matches are offered in the sportsbook at different times, the choice set should not change over time as the same lotteries can be constructed at any point in time from the available matches. The reason is that at any point in time there exists a multitude of simple bets spanning the range of win probabilities, and these can be combined into various bet types to produce the same set of payoff risk profiles. To demonstrate this, in Figure 3 we show that, for two arbitrarily chosen days, it is possible to use simple bets from just a single sport to form binary lotteries whose win probabilities cover the entire range of possible values.<sup>12</sup>

### 2.3 Informational content of lottery choices

Here, we discuss the features of the data that provide information about the model parameters. First, we are able to separate risk preferences from the probability of having the opportunity to bet because, while both affect betting frequency, only preferences affect choices conditional on placing a bet. For example, two individuals who have different probabilities of having the opportunity to bet but the same risk preferences will exhibit different betting frequencies but similar choices conditional on placing a bet. Second, we are able to separate risk aversion from loss aversion by observing individual attitude to lotteries with small vs. large stakes. Risk aversion affects the utility's global concavity while loss aversion primarily affects its concavity at the reference point. Thus, for a given aversion to small-stakes gambles, the more averse an individual is to large-stakes gambles, the more likely he is to be risk rather than loss averse (Rabin, 2000). Third, we are able to separate probability weighting from risk aversion by observing attitude to lotteries with different levels of variance vs. skewness. A risk-seeking individual chooses lotteries with high variance, while one who overweights small probabilities chooses lotteries with high skewness. Finally, an individual whose probability weighting function has an inverse-S shape overweights outcomes at both extremes of the distribution, while one whose probability weighting function is elevated primarily overweights outcomes at one extreme—the left (right) in the case of losses (gains).

## 3 Results

In this section, we present our estimation results. In Section 3.1, we present results on the population proportions of the behavioral types. In Section 3.2, we present results on the distribution of the preference parameters over the entire population. In Section 3.3, we discuss simulation results that show that ignoring heterogeneity within or across behavioral types would yield different results.

<sup>12</sup>Bookmakers do not offer wagers with win probability very close to 1, because they need to make a profit.

### 3.1 Proportions of utility types

Table 3 presents posterior estimates for the population proportions of the four utility types. Our results show that utility curvature alone cannot adequately explain individuals’ choices, and that incorporating the distinctive features of prospect theory—loss aversion and probability weighting—is crucial. To assess the relative importance of these two features, we observe that individuals can be partitioned into two groups: 63% belong to the unrestricted CPT type and therefore exhibit all CPT features, and 36% belong to the type that exhibits probability weighting but not loss aversion. Only 1% of the population belongs to the type that does not exhibit probability weighting. We conclude that, while loss aversion is an important feature of risk attitudes, probability weighting is even more so. This result adds to extant findings by providing richer evidence on the importance of probability distortions in the field. For example, Barseghyan et al. (2013) find that some kind of “generic” probability distortions are important in explaining insurance deductible choices, and Snowberg and Wolfers (2010) show that, with a representative agent, a model with probability weighting alone better fits racetrack betting prices than a model with risk aversion alone. Using a dataset which contains both gains and losses and a model that allows for rich individual heterogeneity, we are able to identify loss aversion and individual preferences, hence to refine and extend these results: Probability weighting is not only an important component of behavior, but furthermore it is substantially more prevalent than loss aversion.

Figure 4 plots a histogram of individuals’ posterior mean of the probability of belonging to the unrestricted type.<sup>13</sup> It shows that about one-third of the individuals are clearly allocated to the unrestricted type, as their posterior probability of belonging to it is close to 1, while the remaining two-thirds are less clearly allocated. Given that we allow for preference heterogeneity both within and across types, it is not surprising that, for some individuals, classification is ambiguous. As we noted previously, the posterior probability that an individual belongs to a type is increasing in the overlap between his likelihood and the distribution of the preference parameters within that type’s sub-population. Thus, the less distinct are the sub-population distributions of the preference parameters, which is the case in the presence of within-type heterogeneity, the less clear is individuals’ classification to types.

### 3.2 Population distribution

Table 4A presents the posterior means and variances of the model parameters, separately for each utility type. Table 4B presents the percentiles and Figure 5 plots the estimated population density for each parameter over the entire population. Finally, Table 5 presents the correlations between the estimated preference parameters.

Reassuringly, our average estimates fall within the range of estimates from lab experiments (see Wakker, 2010 for a review). Specifically, the value function we estimate is moderately concave (convex) over gains (losses), with a median  $\alpha$  of 0.65, and loss aversion is mild, with a median  $\lambda$  of 1.13. The estimates for the curvature (median  $\gamma$  of 0.84) and elevation (median  $\kappa$  of 2.14) of the probability weighting function indicate it is mostly concave, implying significant overweighting of extreme outcomes. Taken together, our average estimates suggest that individuals are averse to risk and exhibit substantial probability weighting

---

<sup>13</sup>Since the estimated proportions of the unrestricted type (63%) and the type with only probability weighting (36%) sum to 99%, the histogram for the latter is a mirror image of the former, while the histograms for the other types are concentrated at zero probability.

so a preference for skewness. Importantly, we find considerable heterogeneity in preferences both across (see Table 3) and within (see Table 4A) utility types, and certainly overall (see Figure 5). Finally, we find a positive relationship between the curvature of the value function and the curvature and elevation of the probability weighting function. This is consistent with the insight in Hsee and Rottenstreich (2004) that common psychological processes drive the sensitivity to deviations from the reference point (captured by  $\alpha$ ) and from impossibility/certainty (captured by  $\gamma$  and  $\kappa$ ). Restating an insight from Qiu and Steiger (2011), the stronger this relationship is, the less of the variation in risk attitudes in a population can EUT alone explain.<sup>14</sup> So our finding further emphasizes the importance of using CPT to explain heterogeneous risk attitudes.<sup>15</sup>

### 3.3 Simulation analysis

Our results are distinct from those that would be obtained by previous approaches that use a mixture of discrete types *or* a hierarchical normal model to estimate heterogeneous CPT preferences (see, e.g., Bruhin, Fehr-Duda and Epper, 2010 and Andrikogiannopoulou and Papakonstantinou, 2020 respectively). To illustrate the dangers of ignoring heterogeneity within or across types, we conduct a simulation analysis in which we generate data for loss aversion  $\lambda$  and probability weighting curvature  $\gamma$  from a distribution that resembles the one we estimate from the real data; we focus on these two parameters to simplify the analysis and facilitate the visual presentation.

In Figure 6a, we present a scatterplot of the simulated data. We see that the simulated data exhibit heterogeneity both within and across two types—one type has mass 35% and exhibits probability weighting but not loss aversion, and the other has mass 65% and exhibits both features.

In Figure 6b, we plot the simulated marginal distribution for loss aversion, together with the estimated marginal from a model that assumes normality hence ignores heterogeneity *across* types. We see that the normal model does not simply shift probability mass from  $\lambda = 1$  to  $\lambda$ s close to 1, which would be innocuous. Rather, it misestimates the probability mass for values as extreme as 0 and 2, and furthermore it misestimates robust summary statistics like the median that are often used to calibrate applied models. This can have very misleading implications about individual and aggregate behavior. For example, the much higher estimated mass of individuals with gain seekingness could incorrectly be thought to explain behaviors such as the reverse disposition effect (as observed by Liu et al., 2010). Furthermore, while with the simulated median  $\lambda$  of 1.15 the implied equity premium would be a mere 0.4%, with the estimated median of 1.4 it would be a substantial 2.8% per year, leading one to incorrectly conclude that loss aversion plays a significant role in explaining this prominent stock market puzzle.<sup>16</sup>

<sup>14</sup>The intuition is as follows. A more curved value function works somewhat like a less curved/elevated probability weighting function: both imply diminished sensitivity to extreme prizes. So if, contrary to our finding, a more curved value function was correlated with less probability weighting, EUT could capture a lot of the variation in risk attitudes despite ignoring probability weighting.

<sup>15</sup>In the Internet Appendix, we also show that our model fits the data quite well. In absolute terms, the stochastic utility component needed to explain the observed choices is small and comparable to values estimated in lab experiments. In relative terms, we find that our model fits the data better than (i) a model in which lotteries are chosen randomly; (ii) a model in which all individuals have the Tversky and Kahneman (1992) parameters; and (iii) a model with no heterogeneity.

<sup>16</sup>The equity premium puzzle is the observation that the mean excess return of equities over bonds (about 6% per year) is inconsistent with standard asset pricing models and plausible risk aversion levels. Our calculations follow those in Benartzi and Thaler (1995).

In Figure 6c, we present results from a model that allows for heterogeneity across types but ignores heterogeneity *within* types. We see that this model estimates two types: one with parameters  $\lambda = 1.04$ ,  $\gamma = 0.91$  and population proportion 65%, and the other with parameters  $\lambda = 2.05$ ,  $\gamma = 0.80$  and population proportion 35%. This estimation misses information in the simulated data, as it collapses them all to two points; e.g., it ignores that a significant minority of individuals (10%) has extreme  $\lambda$ s (below 0.75 or above 2.5). More importantly, it incorrectly concludes that the large majority of individuals (65%) have essentially expected utility theory (EUT) preferences ( $\lambda = 1.04$ ,  $\gamma = 0.91$ ), while in the simulated data all individuals exhibit at least one behavioral feature and only a small minority of 20% are “close” to having EUT preferences ( $\lambda \in (0.9, 1.1)$  and  $\gamma \in (0.9, 1.1)$ ).

## 4 Application to portfolio choice

As we noted above, our average estimates fall within the range of estimates from lab experiments, and so they do not indicate a severe sample selection bias. Despite this, it is important to be careful when extrapolating preferences from *any* group or context to another. In this section, we take our estimates and we apply them to a traditional financial market. Specifically, we study a standard portfolio choice problem, i.e., the optimal decision of an individual who allocates his wealth between the risk-free asset, a single-stock undiversified portfolio, and a well-diversified fund. This analysis enables us to examine how the optimal decisions that we calculate relate to evidence on household portfolio allocations, hence provides a sense of how realistic our estimates are when applied to a choice situation different from the one in which we estimate them.

In the asset-allocation problem we consider, an individual solves

$$\begin{aligned} \max_{\omega_S, \omega_F} U(r_p; \alpha, \lambda, \gamma, \kappa) \\ \text{s.t. } r_p = r_f + \omega_S r_S + \omega_F r_F, \end{aligned} \quad (8)$$

where  $r_f$  is the rate of return of the risk-free asset;  $r_S$  and  $r_F$  are the rates of return (in excess of  $r_f$ ) of the stock and the fund, respectively;  $r_p$  is the rate of return of portfolio  $p$  with portfolio weights  $\omega_S$  and  $\omega_F$  on the stock and the fund, respectively; and  $U$  is the CPT utility functional.<sup>17</sup> Since the distribution of portfolio returns is essentially continuous, the utility functional is an integral rather than a sum, i.e.,

$$U(r_p; \cdot) := \int_{-\infty}^0 u(x; \cdot) \frac{d}{dx} (w(F_p(x); \cdot)) dx + \int_0^{+\infty} u(x; \cdot) \frac{d}{dx} (-w(1 - F_p(x); \cdot)) dx,$$

where  $u$  is the value function and  $w$  the probability weighting function defined in Equations 3 to 5, and  $F_p$  is the distribution function for the returns of portfolio  $p$ . We solve this problem for each of our individuals, using the posterior mean estimates of his preference parameters ( $\alpha_n$ ,  $\lambda_n$ ,  $\gamma_n$ , and  $\kappa_n$ ) and nominal annual returns for  $r_f$ ,  $r_S$ , and  $r_F$ .<sup>18</sup> To solve this problem, we construct a two-dimensional grid spaced at 1%

<sup>17</sup>Since retail investors are unlikely to engage in shorting and/or borrowing, we impose  $\omega_S \geq 0$ ,  $\omega_F \geq 0$ , and  $\omega_S + \omega_F \leq 1$ . Also, since  $U$  is homogeneous in wealth, we can evaluate it on portfolio rates of return rather than on wealth changes as in the more common CPT formulation. See Polkovnichenko (2005) for a similar setup.

<sup>18</sup>Our use of nominal annual returns follows common practice in the literature. Benartzi and Thaler (1995) argue that households consider nominal annual returns in portfolio allocation because they file taxes and receive reports from brokers and mutual funds annually, and nominal returns are more prominent in annual reports.

increments over the range of values for the weights  $\omega_S$  and  $\omega_F$ , and we pick the weights that maximize the objective  $U$  for each individual. We approximate the objective by generating 1 million draws for the returns  $r_f$ ,  $r_S$ , and  $r_F$  from their empirical joint distribution, and numerically calculating the integrals.<sup>19</sup>

We find that 48% of our individuals would optimally invest in equity and, conditional on investing in equity, the proportion of undiversified to total equity investment varies widely across individuals (see Figure 7) with a mean/median of 49%/46% and a standard deviation of 24%. These results are broadly in line with corresponding household data from the U.S. Survey for Consumer Finances (SCF). First, according to the SCF, since 2001 about 50% of all households in the U.S. participate in the stock market each year, which is similar to the 48% participation rate we calculate using our CPT estimates. This finding — that half of our individuals would not participate in a market that offers risky assets with positive expected returns — is not surprising given that we find that individuals exhibit a preference for *skewness*, not *risk*. Second, both in the SCF data and in our analysis, a significant proportion invests simultaneously in both undiversified and well-diversified equity portfolios: In the 2001 SCF, the proportion of direct to total equity holdings has a median value of 40% across households, which is close to the median value of 46% that we calculate. This behavior—which is puzzling with expected utility preferences—is optimal for a significant proportion of our individuals due to the interaction of probability weighting which makes the undiversified portfolios (with their low probability of very large gains) attractive, and loss aversion which makes the diversified portfolios (with their low probability of big losses coupled with decent average gains) attractive. In Section IA.5 of the Internet Appendix, we analyze how each preference parameter affects the optimal weights on the stock and the fund in order to see how our estimated distribution of preference parameters translates to the distribution of optimal portfolios we present in Figure 7.

Overall, we conclude that our preference parameter estimates yield portfolios similar to those observed in the data. This indicates that our estimates are sensible and could be relevant more generally.

## 5 Extensions and robustness checks

Here, we present extensions and robustness checks to our main analysis: (i) we consider alternative specifications for beliefs, (ii) we impose a budget constraint, (iii) we test for the prevalence of betting skill, and (iv) we conduct a prior sensitivity analysis.

### 5.1 Alternative beliefs

In our main analysis, we follow the vast majority of the literature on preference estimation and we assume beliefs are rational, approximating the outcomes' true probabilities with those implied by their quoted odds. Here, we discuss two alternative belief specifications: one that uses a different approximation of

---

<sup>19</sup>Specifically, for the period 1975–2011, we obtain data on the risk-free rate (the one-month T-Bill) and on equity returns from the Center for Research in Securities Prices (CRSP), and on U.S. equity mutual fund returns (net of fees, expenses, and transaction costs) from the CRSP U.S. Mutual Fund Database. Then, for each draw, we randomly pick a month  $t$  in this period, we randomly pick a stock among all stocks and a fund among all funds active in month  $t$ , and we calculate the returns for the twelve months starting with month  $t$ . The mean/median/standard deviation of our draws are 5.51%/5.23%/3.26% for  $r_f$ , 12.07%/4.69%/79.60% for  $r_S$ , and 11.49%/12.54%/24.52% for  $r_F$ . For details, see Section IA.5 of the Internet Appendix.

rational beliefs and one that allows for heterogeneous subjective beliefs.

First, we approximate the outcomes’ true probabilities using the win frequencies of outcomes with similar odds. In Figure 8a, we compare the posteriors from this analysis with those from the baseline, and we see that they are quite similar. This is not surprising, given that the two sets of probabilities are similar (see Figure 2).

Second, we allow for the possibility that observed choices are at least partially motivated by heterogeneous optimistic subjective beliefs. Letting  $1-p$  be the implied win probability for a selected outcome, we let  $1-\zeta p$  for  $\zeta \in [\underline{\zeta}, 1]$  be its perceived win probability, that is, the individual perceives the probabilities of outcomes that he has backed to be higher than the probabilities of these outcomes implied by the prices quoted by the bookmaker. This analysis deviates from the usual assumption of rationality and is in the general spirit of that in Gandhi and Serrano-Padial (2014) and Ericson et al. (2021), who study heterogeneous beliefs. Interestingly the latter find that, if they allow for heterogeneous beliefs, the insurance choices they observe are consistent with homogeneity in risk preferences. Contrary to them, as we can see in Figure 8b, we still find a lot of heterogeneity in all risk preference parameters. Specifically, in Figure 8b, we plot the posteriors for two values of  $\underline{\zeta}$ , hence for two deviation magnitudes: across all lotteries, the small/large deviations increase the perceived probability of winning the maximum prize by 1.4%/3.3% on average, and increase the lottery’s expected value by 2.2%/4.5%. We find that these deviations have little effect on our results. Intuitively this is because, even if an individual bets on a match because he (believes he) has superior information on it, he still needs to choose among the multitude of wagers that involve this match, and he also has the option to combine this wager with wagers on other matches. For example, if an individual considers backing Barcelona in its soccer match against Real Madrid because he believes that Barcelona’s form is under-valued relative to that of Real Madrid, he still needs to choose what kind of bet to place, e.g., to back Barcelona or to back Barcelona to win by a large margin. Furthermore, the individual has the option to combine a wager on this event with wagers on other events. These choices lead to lotteries involving very different risk profiles, so the observed variation in choices cannot be explained solely by variation in beliefs, but rather is caused by (and so informs us about) variation in individuals’ preferences.

## 5.2 Budget-constrained choice set

In our baseline analysis, we construct the choice set by drawing the stake and other lottery characteristics from their empirical distributions in our data. Here, we account for the possibility that individuals may face a budget constraint. Specifically, we reconstruct the choice set by constraining the maximum possible stake to be €500. This is a reasonable bound given that (i) the average annual income of the individuals under study proxied by matching their gender, age, and zip code with the relevant census data is around 24,000 euros, and (ii) it has been found that even the least liquid households hold considerable liquidity (cash plus credit capacity) of approximately 60% of their monthly income (Olafsson and Pagel, 2018). In Figure 9, we compare the posteriors from this analysis with those from the baseline analysis, and we see that they are very similar. The intuition for this is that individuals rarely choose lotteries with intermediate stakes, e.g., between €100 and €500, therefore the inclusion or exclusion of lotteries with even higher stakes does not reveal substantially different information regarding their preferences.

### 5.3 Testing for skill in picking bets

Here, we examine how many individuals exhibit ‘skill’ in picking bets, hence their choices might not (solely) reflect their risk preferences. Skilled individuals should be able to exploit mispriced bets to systematically earn returns in excess of the expected return implied by the prices. To detect skill, we carry out individual-level tests to examine whether any individual’s average excess return significantly exceeds 0; since we perform multiple hypothesis tests at the individual level, we calculate bootstrap  $p$ -values and use the Sidak (1967) correction for significance levels. We find no skilled individuals in our sample.

### 5.4 Prior sensitivity analysis

Here, we focus on the sensitivity of the *posterior* predictive densities of the parameters on their *prior* predictive densities, where the prior (posterior) predictive densities are the distributions of an individual’s parameters before (after) observing our data for all individuals. Figure 10 shows that a drastic change in the prior predictive density — from the baseline that is flat almost everywhere, except for values close to the endpoints for bounded distributions, to an alternative that is concentrated over a smaller range of values away from the boundaries — has almost no effect on the posterior predictive density, for all parameters.

In the Internet Appendix, we present additional results in which we vary the marginal priors of the population parameters, rather than the prior predictive densities which compound all the marginals. Specifically, we show that our results are robust to varying by orders of magnitude the scale  $\underline{\Delta}$  of the prior distribution for  $V_{\rho}^q$ , which shows that the heterogeneity we estimate is not driven by the priors. We also show that varying the priors for the utility type proportions  $h$  does not affect their posteriors, indicating that the priors do not drive our finding that a very substantial 36% of the individuals do not exhibit loss aversion.

## 6 Conclusion

The paper makes two separate contributions. First, we develop a hierarchical mixture model that allows individual preferences to vary both within and across behavioral types, allowing us to estimate the importance and prevalence of behavioral features in the population. Our model and estimation technique are generally applicable in a variety of experimental and field settings in which substantial heterogeneity is likely to exist both within and across groups. For example, it can be used to estimate the proportions of firms with pro- or anti-social practices, of managers with positive versus negative skill, or of individuals with exponential versus hyperbolic time preferences. Second, our methodology yields novel insights when applied to a rich panel of behavior in a sports betting market. We find that prospect theory should not be viewed as a monolithic theory with a set of features—loss aversion and probability weighting—that all individuals exhibit, but rather reality is more nuanced: utility curvature alone does not explain observed choices and, while two-thirds of individuals exhibit loss aversion, all exhibit probability weighting. These findings could guide the use of prospect theory in applied models. For example, our estimated proportions could serve as valuable inputs for theoretical models that involve different types of behavioral agents, and furthermore allowing for such heterogeneity in models may be crucial to understanding observed patterns in behavior at the individual and aggregate levels (e.g., portfolio allocation and asset prices).

## References

- Abdellaoui, M.** 2000. "Parameter-free elicitation of utility and probability weighting functions." *Management Science*, 46(11): 1497–1512.
- Abdellaoui, M, H Bleichrodt, and C Paraschiv.** 2007. "Loss aversion under prospect theory: A parameter-free measurement." *Management Science*, 53(10): 1659–1674.
- Andrikogiannopoulou, A, and F Papakonstantinou.** 2020. "History-dependent risk preferences: Evidence from individual choices and implications for the disposition effect." *Review of Financial Studies*, 33(8): 3674–3718.
- Apesteguia, J, and M. A Ballester.** 2018. "Monotone stochastic choice models: The case of risk and time preferences." *Journal of Political Economy*, 126(1): 74–106.
- Baillon, A, H Bleichrodt, and V Spinu.** 2020. "Searching for the reference point." *Management Science*, 66(1): 93–112.
- Barberis, N.** 2012. "A model of casino gambling." *Management Science*, 58(1): 35–51.
- Barberis, N.** 2013. "Thirty years of prospect theory in economics: A review and assessment." *Journal of Economic Perspectives*, 27(1): 173–96.
- Barseghyan, L, F Molinari, T O'Donoghue, and J. C Teitelbaum.** 2013. "The nature of risk preferences: Evidence from insurance choices." *American Economic Review*, 103(6): 2499–2529.
- Barseghyan, L, F Molinari, T O'Donoghue, and J. C Teitelbaum.** 2018. "Estimating risk preferences in the field." *Journal of Economic Literature*, 56(2): 501–564.
- 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*, ed. J. Volmuller and R. Hamerslag. 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(1): 73–92.
- Brownstone, D, D. S Bunch, and K Train.** 2000. "Joint mixed logit models of stated and revealed preferences for alternative-fuel vehicles." *Transportation Research Part B: Methodological*, 34(5): 315–338.
- Bruhin, A, H Fehr-Duda, and T Epper.** 2010. "Risk and rationality: Uncovering heterogeneity in probability distortion." *Econometrica*, 78(4): 1375–1412.
- Cohen, A, and L Einav.** 2007. "Estimating risk preferences from deductible choice." *American Economic Review*, 97(3): 745–788.
- Conte, A, J. D Hey, and P. G Moffatt.** 2011. "Mixture models of choice under risk." *Journal of Econometrics*, 162(1): 79–88.
- Ericson, K. M, P Kircher, J Spinnewijn, and A Starc.** 2021. "Inferring risk perceptions and preferences using choice from insurance menus: Theory and evidence." *The Economic Journal*, 131(634): 713–744.
- Fasman, J.** 2010. "Shuffle up and deal." *Economist*, July 8. <http://tinyurl.com/gstoqu2>, Last Accessed: 20 October, 2021.
- Gandhi, A, and R Serrano-Padial.** 2014. "Does belief heterogeneity explain asset prices: The case of the longshot bias." *The Review of Economic Studies*, 82(1): 156–186.

- von Gaudecker, H, A van Soest, and E Wengstrom.** 2011. "Heterogeneity in risky choice behavior in a broad population." *American Economic Review*, 101(2): 664–694.
- Geweke, J.** 1999. "Using simulation methods for Bayesian econometric models: Inference, development, and communication." *Econometric Reviews*, 18(1): 1–73.
- Goldstein, W. M, and H. J Einhorn.** 1987. "Expression theory and the preference reversal phenomena." *Psychological Review*, 94(2): 236–254.
- H2 Gambling Capital.** 2013. "There's nothing virtual about the opportunity in real-money gambling." H2 Gambling Capital & Odobo, Gibraltar.
- Hastings, W.** 1970. "Monte Carlo sampling methods using Markov chains and their applications." *Biometrika*, 57(1): 97–109.
- Hsee, C. K, and Y Rottenstreich.** 2004. "Music, pandas, and muggers: On the affective psychology of value." *Journal of Experimental Psychology: General*, 133(1): 23–30.
- Jullien, B, and B Salanie.** 2000. "Estimating preferences under risk: The case of racetrack bettors." *Journal of Political Economy*, 108(3): 503–530.
- Kahneman, D, and A Tversky.** 1979. "Prospect theory: An analysis of decisions under risk." *Econometrica*, 47(2): 263–292.
- Kobberling, V, and P. P Wakker.** 2005. "An index of loss aversion." *Journal of Economic Theory*, 122(1): 119–131.
- Liu, Y.-J, C.-L Tsai, M.-C Wang, and N Zhu.** 2010. "Prior consequences and subsequent risk taking: New field evidence from the Taiwan Futures Exchange." *Management Science*, 56(4): 606–620.
- Manski, C. F.** 2004. "Measuring expectations." *Econometrica*, 72(5): 1329–1376.
- Metropolis, N, A Rosenbluth, M Rosenbluth, A Teller, and E Teller.** 1953. "Equation of state calculations by fast computing machines." *Journal of Chemical Physics*, 21(6): 1087–1092.
- Nilsson, H, J Rieskamp, and E.-J Wagenmakers.** 2011. "Hierarchical Bayesian parameter estimation for cumulative prospect theory." *Journal of Mathematical Psychology*, 55(1): 84–93.
- O'Donoghue, T, and C Sprenger.** 2018. "Reference-dependent preferences." In *Handbook of behavioral economics*, Vol. 1, ed. Douglas Bernheim, Stefano DellaVigna and David Laibson, Chapter 1, 1–77. Amsterdam, the Netherlands:Elsevier.
- Olafsson, A, and M Pagel.** 2018. "The liquid hand-to-mouth: Evidence from personal finance management software." *The Review of Financial Studies*, 31(11): 4398–4446.
- Polkovnichenko, V.** 2005. "Household portfolio diversification: A case for rank-dependent preferences." *Review of Financial Studies*, 18(4): 1467–1502.
- Polkovnichenko, V, and F Zhao.** 2013. "Probability weighting functions implied in options prices." *Journal of Financial Economics*, 107(3): 580–609.
- Pope, P. F, and D. A Peel.** 1989. "Information, prices and efficiency in a fixed-odds betting market." *Economica*, 56(223): 323–41.
- Post, T, M. J van den Assem, G Baltussen, and R. H Thaler.** 2008. "Deal or no deal? Decision making under risk in a large-payoff game show." *American Economic Review*, 98(1): 38–71.

- Qiu, J, and E Steiger.** 2011. "Understanding the two components of risk attitudes: An experimental analysis." *Management Science*, 57(1): 193–199.
- Rabin, M.** 2000. "Risk aversion and expected utility theory: A calibration theorem." *Econometrica*, 68(5): 1281–1292.
- Robert, C. P.** 2007. *The Bayesian choice: From decision-theoretic foundations to computational implementation*. New York:Springer Science & Business Media.
- Sidak, Z.** 1967. "Rectangular confidence regions for the means of multivariate normal distributions." *Journal of the American Statistical Association*, 62(318): 626–633.
- Snowberg, E, and J Wolfers.** 2010. "Explaining the favorite-longshot bias: Is it risk-love or misperceptions?" *Journal of Political Economy*, 118(4): 723–746.
- Tversky, A, and D Kahneman.** 1992. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and Uncertainty*, 5(4): 297–323.
- Wakker, P. P.** 2010. *Prospect theory: For risk and ambiguity*. Cambridge University Press.
- Woodland, L. M, and B. M Woodland.** 1994. "Market efficiency and the favorite-longshot bias: The baseball betting market." *Journal of Finance*, 49(1): 269–79.

Table 1: Summary Statistics for Individual Characteristics and Betting Activity

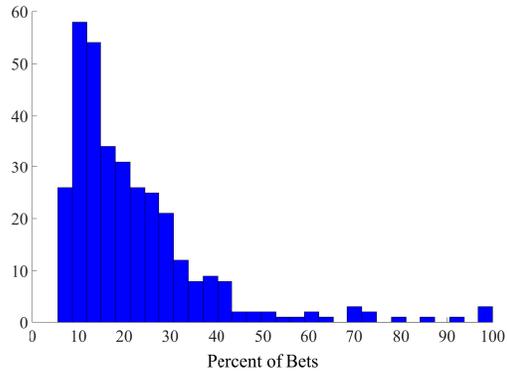
In Panel A, Female is a dummy indicating gender. Age is in years. In Panel B, Event types / Leagues / Teams per individual is the number of different event types / leagues / teams on which each individual has bet. A bet day for an individual is a day on which he has placed a bet. Bet days per year per individual measures betting frequency. Events per bet day is the number of events on which an individual bets in a bet day.

Panel A: Individual Characteristics

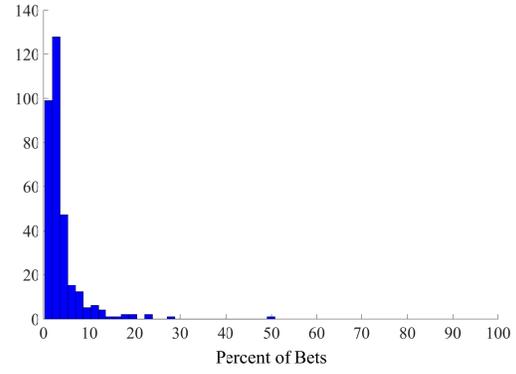
	Mean	Median	Std. Dev.	Min	Max
Female	0.07	0	0.25	0	1
Age	32.85	31	9.63	18	67

Panel B: Bet Characteristics

	Mean	Median	Std. Dev.	Min	Max
Event types per individual	23.11	15	25.49	1	235
Leagues per individual	34.51	32	20.64	1	109
Teams per individual	182.61	122	174.81	4	976
Bet days per individual	35.02	21	42.03	5	380
Bet days per year, per individual	64.76	42	63.93	1	327
Events per bet day	5.44	4	5.16	1	66



(a) Bets on the most-frequently-bet league.



(b) Bets backing the most-frequently-bet team.

Figure 1: Histograms, across individuals, of betting activity involving specific leagues (Panel *a*) and specific teams (Panel *b*).

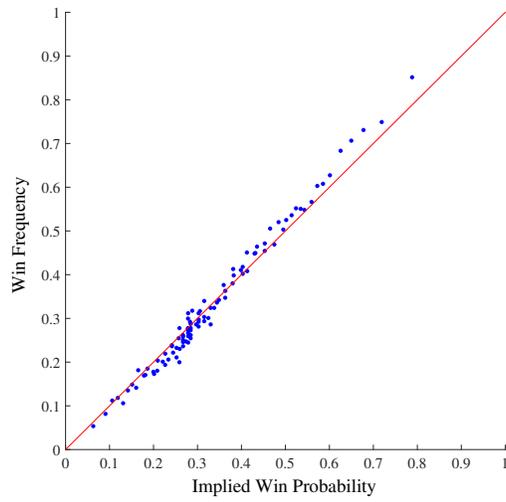


Figure 2: Comparison of the win probabilities implied by the quoted prices with the historical win frequencies of outcomes with similar quoted prices, for the 100 percentile price groups. The red solid straight line plots the 45° line.

Table 2: Summary Statistics

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

	Percentiles						
	1 <sup>st</sup>	5 <sup>th</sup>	25 <sup>th</sup>	50 <sup>th</sup>	75 <sup>th</sup>	95 <sup>th</sup>	99 <sup>th</sup>
Number of Prizes	2	2	2	2	6	61	354
Mean	-43.39	-17.54	-4.96	-2.05	-0.77	-0.13	-0.03
Standard Deviation	0.31	1.43	8.18	20.17	45.03	153.91	411.11
Skewness	-1.18	-0.32	0.72	2.49	5.94	33.23	180.94
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	$3 \cdot 10^3$	$2 \cdot 10^4$

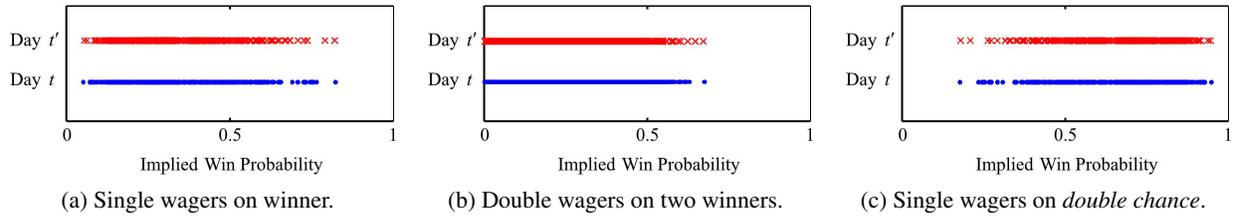


Figure 3: Illustration of the wide range of win probabilities available on the sportsbook on two arbitrary days, when placing simple bets (which can be combined into complex bets to generate a wide array of risk profiles). Specifically, we plot the implied win probabilities for two arbitrarily chosen days—in blue dots for one day ( $t$ ) and in red x-marks for the other day ( $t'$ )—for simple binary bets that can be placed on soccer matches in major leagues. Panel (a) plots the win probabilities for individual wagers on the winner, i.e., bets that back the home team, the draw, or the away team. Panel (b) plots the win probabilities for bets that combine two individual wagers on the winners of two soccer matches. Panel (c) plots the win probabilities for ‘double chance’ individual wagers, i.e., bets that back the home or draw, the home or away, or the draw or away outcome.

Table 3: Utility Type Proportions

Summary statistics of the posterior distribution for the population proportions of the four behavioral types: one exhibits neither loss aversion nor probability weighting ( $\lambda = \gamma = \kappa = 1$ ), one does not exhibit loss aversion ( $\lambda = 1$ ), one does not exhibit probability weighting ( $\gamma = \kappa = 1$ ), and one exhibits all CPT features. The 95% Highest Posterior Density Interval (HPDI) for a parameter is the smallest interval such that the posterior probability that the parameter is in this interval is 0.95.

	Mean	Median	Std.Dev.	95% HPDI
Type with $\lambda = \gamma = \kappa = 1$	0.00	0.00	0.00	[0.00, 0.02]
Type with $\lambda = 1$	0.36	0.36	0.04	[0.28, 0.44]
Type with $\gamma = \kappa = 1$	0.01	0.01	0.01	[0.00, 0.03]
Unrestricted Type	0.63	0.63	0.04	[0.54, 0.71]

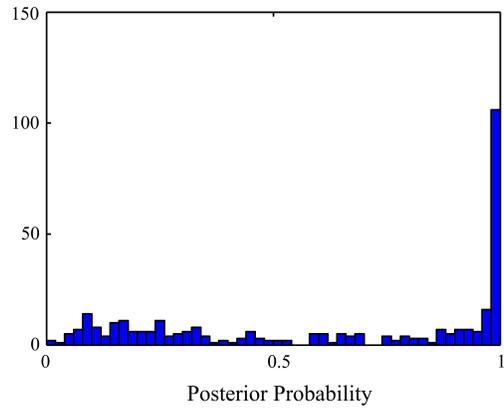


Figure 4: Histogram of individuals' posterior probability of belonging to the unrestricted full-featured type.

Table 4: Posterior Estimates of the Model Parameters' Population Distribution

Panel A shows summary statistics of the posterior estimates for the population mean and variance of  $\pi$ ,  $\tau$ , and  $\rho$  for each of the four utility types estimated using our mixture model;  $\pi$  and  $\tau$  are common across all types, so their means and variances take the same values for all types. Panel B shows the posterior means of the percentiles of  $\pi$ ,  $\tau$ , and  $\rho$ , calculated over the entire population.

Panel A: Means and Variances

	$\lambda = \gamma = \kappa = 1$ Type		$\lambda = 1$ Type		$\gamma = \kappa = 1$ Type		Unrestricted Type		
	Mean	Std.Dev.	Mean	Std.Dev.	Mean	Std.Dev.	Mean	Std.Dev.	
<i>Means</i>									
$\pi$	0.25	0.01	0.25	0.01	0.25	0.01	0.25	0.01	
$\tau$	1.93	0.10	1.93	0.10	1.93	0.10	1.93	0.10	
$\alpha$	0.88	0.88	0.64	0.02	0.29	0.06	0.67	0.01	
$\lambda$	1	0	1	0	3.73	1.32	1.68	0.08	
$\gamma$	1	0	0.88	0.03	1	0	0.81	0.01	
$\kappa$	1	0	1.91	0.10	1	0	3.19	0.21	
<i>Variances</i>									
$\pi$	0.05	0.00	0.05	0.00	0.05	0.00	0.05	0.00	
$\tau$	2.12	0.26	2.12	0.26	2.12	0.26	2.12	0.26	
$\alpha$	0.01	0.13	0.01	0.00	0.00	0.00	0.02	0.00	
$\lambda$	0	0	0	0	0.35	0.17	0.46	0.08	
$\gamma$	0	0	0.01	0.00	0	0	0.02	0.00	
$\kappa$	0	0	0.17	0.06	0	0	3.58	0.54	

Panel B: Percentiles

	1 <sup>st</sup>	5 <sup>th</sup>	25 <sup>th</sup>	50 <sup>th</sup>	75 <sup>th</sup>	95 <sup>th</sup>	99 <sup>th</sup>
$\pi$	0.01	0.02	0.07	0.18	0.38	0.72	0.88
$\tau$	0.16	0.33	0.84	1.53	2.65	4.94	6.64
$\alpha$	0.33	0.45	0.57	0.65	0.74	0.90	1.03
$\lambda$	0.61	0.85	1.00	1.13	1.77	2.81	3.77
$\gamma$	0.55	0.64	0.76	0.84	0.92	1.04	1.13
$\kappa$	0.48	0.89	1.61	2.14	3.38	6.28	7.89

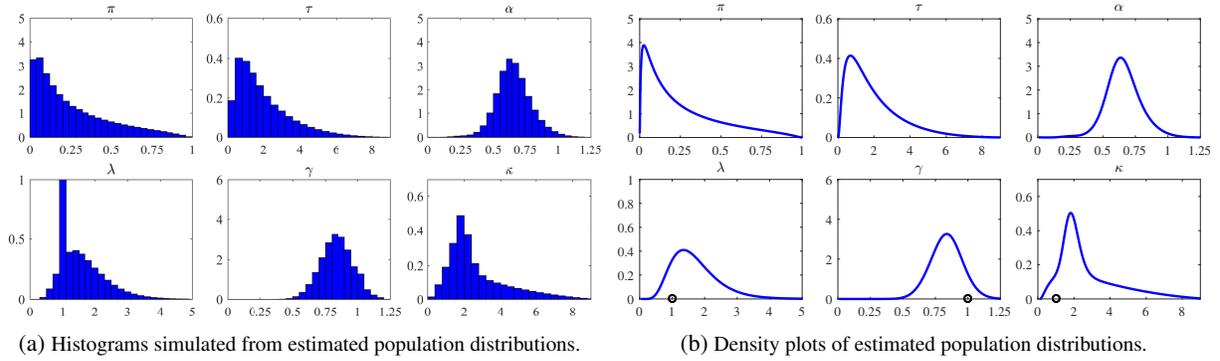


Figure 5: Estimated population distributions of the probability  $\pi$  of having the opportunity to place a bet on any given day, the inverse scale  $\tau$  of the stochastic component of utility, the measures  $\alpha$  and  $\lambda$  of curvature and loss aversion of the value function, and the measures  $\gamma$  and  $\kappa$  of curvature and elevation of the probability weighting function. In Panel (a), we plot histograms of values simulated from the estimated distribution of each parameter. Bar heights are normalized so that the area of each bar represents the probability of the corresponding interval. In Panel (b), we plot the estimated population densities. The circles at  $\lambda = 1$  and at  $\gamma = 1, \kappa = 1$  represent point masses that correspond to the proportions of individuals with no loss aversion and with no probability weighting, respectively. The proportion with no loss aversion is 36% and the proportion with no probability weighting is 1%.

Table 5: Population Correlations

Posterior means and standard deviations (in parentheses) of the correlations between the preference parameters.

	$\lambda$	$\gamma$	$\kappa$
$\alpha$	-0.10 (0.05)	0.33 (0.07)	-0.40 (0.05)
$\lambda$		-0.36 (0.06)	0.63 (0.06)
$\gamma$			-0.15 (0.07)

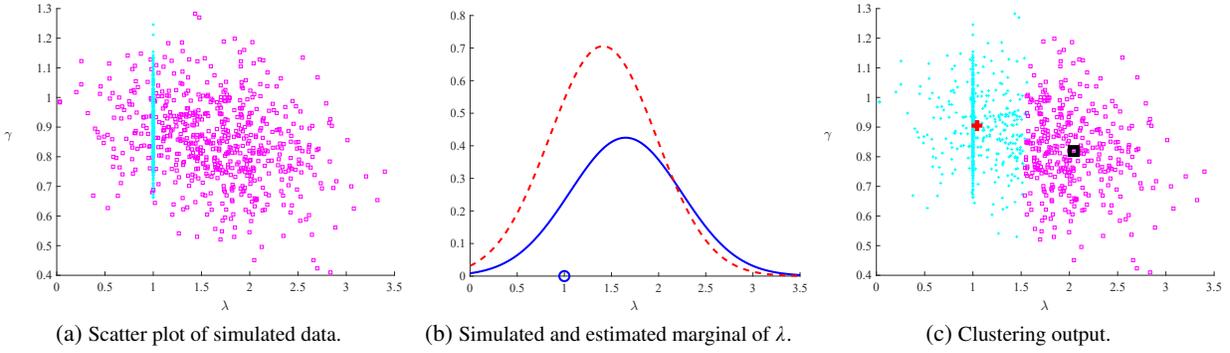


Figure 6: Illustration of the deficiencies of ignoring heterogeneity within or across types. In Panel (a), we present a scatter plot of data simulated from (essentially) the  $(\lambda, \gamma)$  bivariate distribution we estimate from our sports betting data; specifically we draw from a mixture of two components, one (plotted with magenta squares) has weight 0.65 and  $(\lambda, \gamma)' \sim \mathcal{N}\left(\begin{pmatrix} 1.65 \\ 0.85 \end{pmatrix}, \begin{pmatrix} 0.35 & -0.03 \\ -0.03 & 0.02 \end{pmatrix}\right)$  and the other (plotted with cyan crosses) has weight 0.35 and  $\lambda = 1, \gamma \sim \mathcal{N}(0.95, 0.01)$ . In Panel (b), we plot the marginal density of  $\lambda$  for the simulated data (in solid blue) and for the distribution estimated from a model that assumes normality and ignores heterogeneity across types (in dashed red); the circle at  $\lambda = 1$  represents the point mass with weight 35%. In Panel (c), we present the clustering output from a model that ignores heterogeneity within types; one cluster contains the points plotted with cyan crosses and has center at  $(\lambda = 1.04, \gamma = 0.91)$ , indicated with a large red cross, and the other cluster contains the points plotted with magenta squares and has center at  $(\lambda = 2.05, \gamma = 0.80)$ , indicated with a large black square.

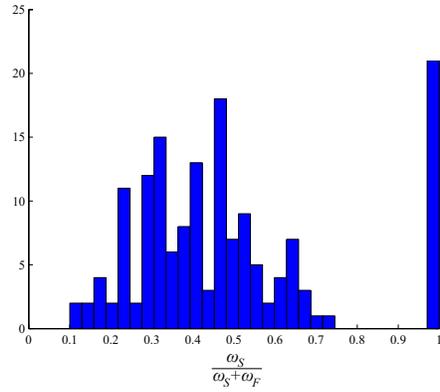


Figure 7: Histogram of the proportion of undiversified to total equity investment across individuals, conditional on investing in equity, as implied by the individual-level preferences we estimate.  $\omega_S$  is the portfolio weight on the stock and  $\omega_F$  is the portfolio weight on the fund, so  $\omega_S$  is the undiversified and  $\omega_S + \omega_F$  is the total equity investment weight.

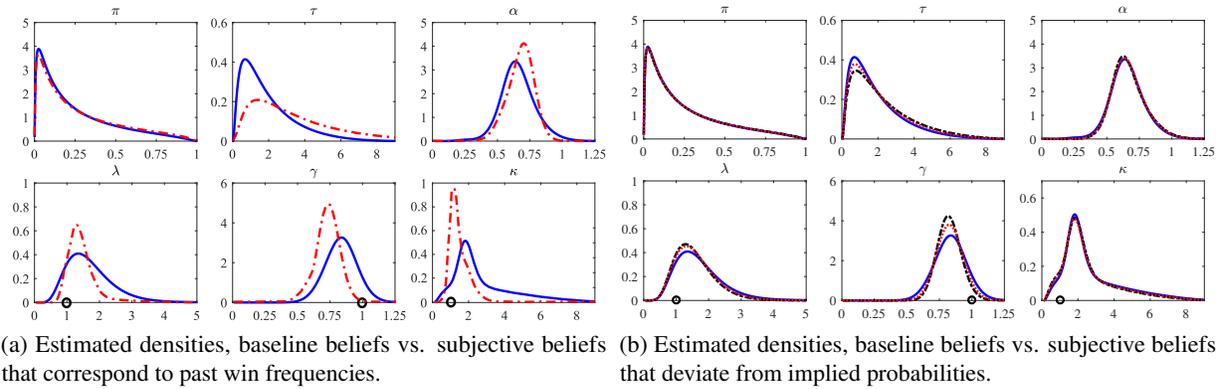


Figure 8: Sensitivity analysis to alternative assumptions about beliefs. In Panel (a), we plot using red dash-dotted lines the estimated densities under the assumption that the subjective probabilities of bet outcomes equal the win frequencies of past outcomes with similar prices. In Panel (b), we plot using red dotted and black dash-dotted lines, respectively, the estimated densities under the assumption that subjective beliefs exhibit small and large deviations from the rational beliefs, as approximated in our baseline analysis by the probabilities implied by the odds. In both panels, we plot in blue solid lines the estimated densities from our baseline analysis. In both panels, the circles at  $\lambda = 1$  and at  $\gamma = 1, \kappa = 1$  represent point masses that correspond to the proportions of individuals with no loss aversion and with no probability weighting, respectively. The proportion with no loss aversion is 36% for the baseline beliefs, 48% for the alternative beliefs in Panel (a), and 35% for the alternative beliefs in Panel (b). The proportion with no probability weighting is 1% in all cases.

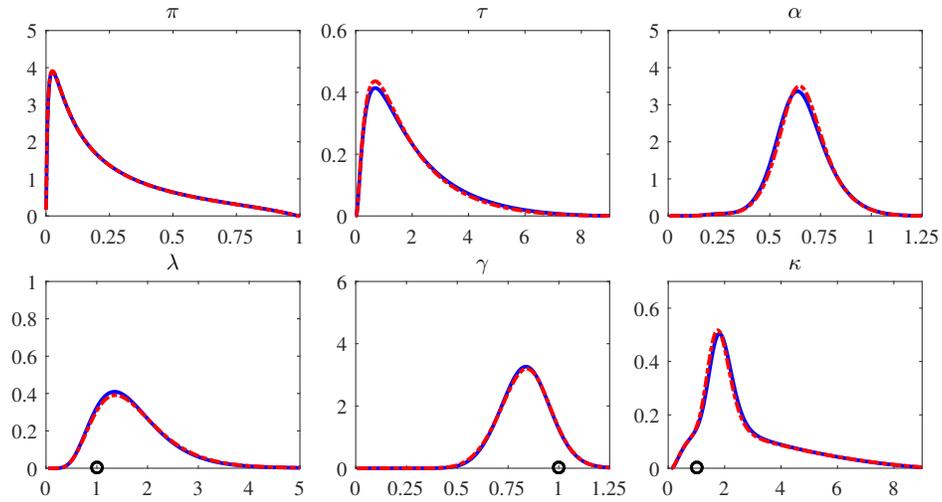


Figure 9: Estimated population densities from our baseline model (blue solid line) and a model in which the maximum possible stake is constrained to be €500 (red dash-dotted line). The proportion with no loss aversion is 36% for the baseline and 37% for the alternative model, and the proportion with no probability weighting is 1% for both models.

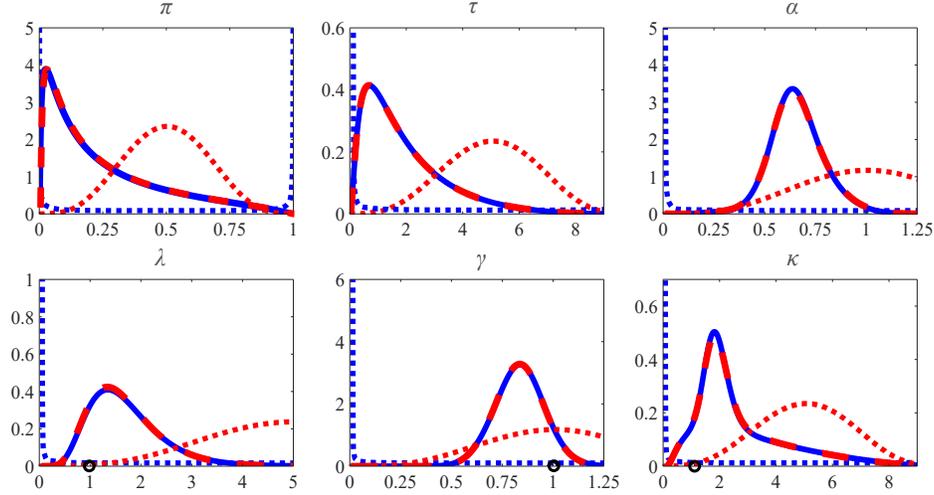


Figure 10: Prior sensitivity analysis. We plot the prior and posterior predictive densities for the baseline prior specification ( $\underline{\kappa} = 0$ ,  $\underline{K} = 100$ ,  $\underline{\lambda} = 6$ ,  $\underline{\Lambda} = I$ ) and for an alternative ( $\underline{\kappa} = 0$ ,  $\underline{K} = 0.5$ ,  $\underline{\lambda} = 6$ ,  $\underline{\Lambda} = I$ ). The prior predictive densities for the baseline (alternative) prior specification are plotted in dotted blue (red) lines; the baseline prior predictive density is flat almost everywhere, except for values close to the endpoints for bounded distributions, while the alternative is concentrated over a smaller range of values away from the boundaries. The proportion with no loss aversion (no probability weighting) is 50% (50%) for both priors. The posterior predictive densities corresponding to the baseline (alternative) prior specification are plotted in solid blue (dashed red) lines, and they almost coincide. The proportion with no loss aversion (no probability weighting) is 36% (1%) for both posteriors. The circles at  $\lambda = 1$  and at  $\gamma = 1$ ,  $\kappa = 1$  represent point masses that correspond to the proportions of individuals with no loss aversion and with no probability weighting, respectively.

## Appendix

### A Choice set generation

To generate the choice set, we follow a procedure that resembles the one people employ when they place bets, i.e., we draw the number of bets in the lottery, the type and number of wagers in each bet, the odds associated with each wager, and the amount staked. We draw these characteristics from their empirical distributions:

1. For the number of bets that comprise a lottery, we draw from a negative binomial that we fit to the number of bets per day lottery across all day lotteries we observe.
2. We draw the general type of each bet (permutation or full-cover) from a Bernoulli with success probability equal to the proportion of permutation bets in the data. The specific bet type depends on the number of bets in the lottery, so we draw the former conditional on the latter. Specifically, for each value of the number of bets in a lottery, we fit a negative binomial to all bet types in lotteries with this number of bets in our data, and then draw from the fitted distributions conditional on the number of bets in the lottery.
3. Furthermore, the number of wagers in a bet is related to the bet type. First, the bet type may impose the precise or minimum number of wagers. Second, the distribution of the number of wagers depends on the bet type. So, for each bet type, we fit a negative binomial with appropriate support to the number of wagers in all bets of that type, and then we draw from the fitted distributions conditional on bet type.
4. The amount staked in each bet also depends on the bet type. We fit to the amounts staked a log-normal whose mean is linear in the type of accumulator bets that the bet type involves. Then, we draw from these fitted distributions the amount staked in a bet conditional on its bet type.
5. For the wager odds, we draw the commission and the implied win probability, and calculate the corresponding odds. The commission is drawn from a Johnson distribution fitted to the commissions in all wagers in our data. The implied win probability is drawn from a uniform ranging from the minimum we observe in our data to the maximum possible given the commission.

We randomly draw 20,000 lotteries and group them into 100 clusters using the Ward hierarchical agglomeration method; it initially places each lottery in its own cluster and in each step combines two clusters so that the mean distance between a lottery and the corresponding cluster center is minimized, where the distance between densities  $f_X, f_Y$  is the Wasserstein distance ( $\inf_{f_{XY}} E [\|X - Y\|]$  with  $f_{XY}$  any joint for  $X, Y$  with marginals  $f_X, f_Y$ ). Then, we choose the most representative lottery in each cluster, i.e., the one with the smallest mean distance to other cluster members. These lotteries reasonably span the randomly drawn lotteries.

### B Markov chain Monte Carlo algorithm

We present our MCMC algorithm for the case where all elements of  $\beta := (\pi, \tau, \rho)'$  have type-specific distributions. We use a Gibbs sampler to draw from the joint density

$$p\left(\mathbf{h}, \{e_n\}, \left\{\mu_{\beta}^q, \mathbf{V}_{\beta}^q\right\}, \{\beta_n\} \mid \{x_n, y_n\}\right),$$

which is proportional to the likelihood in Equation 1 of the main text times the joint prior

$$\left[ \prod_n p\left(\beta_n \mid e_n, \left\{\mu_{\beta}^q\right\}, \left\{\mathbf{V}_{\beta}^q\right\}\right) \cdot p\left(e_n \mid \mathbf{h}\right) \right] \cdot \left[ \prod_q p\left(\mu_{\beta}^q, \mathbf{V}_{\beta}^q\right) \right] \cdot p(\mathbf{h}).$$

The Gibbs sampler consists of one block each for  $\mathbf{h}$ , for  $\{\boldsymbol{\mu}_{\tilde{\beta}}^q, \mathbf{V}_{\tilde{\beta}}^q\}$ , and for  $\{\beta_n\}, \{e_n\}$ .

We start by drawing from the priors. These starting values do not matter if the whole posterior is explored, and indeed results are robust across seeds. Given values in iteration  $k$ , in iteration  $k+1$  we draw from conditionals  $p(\mathbf{h}|\{e_n\}^{(k)})$ ,  $p(\{\boldsymbol{\mu}_{\tilde{\beta}}^q, \mathbf{V}_{\tilde{\beta}}^q\}|\{\beta_n\}^{(k)}, \{e_n\}^{(k)})$ , and  $p(\{\beta_n\}, \{e_n\}|\{\mathbf{x}_n, \mathbf{y}_n\}, \{\boldsymbol{\mu}_{\tilde{\beta}}^q\}^{(k+1)}, \{\mathbf{V}_{\tilde{\beta}}^q\}^{(k+1)}, \mathbf{h}^{(k+1)})$ , which follow from Section 1 of the main text. Specifically, suppressing iteration superscripts, we have:

- Letting  $\bar{\mathbf{h}} = \underline{\mathbf{h}} + \sum_{n=1}^N e_n$ , the conditional posterior for  $\mathbf{h}$  is Dirichlet,  $\mathbf{h}|\{e_n\} \sim D(\bar{\mathbf{h}})$ .
- The conditional posterior for  $(\boldsymbol{\mu}_{\tilde{\beta}}^q, \mathbf{V}_{\tilde{\beta}}^q)$  is

$$\boldsymbol{\mu}_{\tilde{\beta}}^q|\{e_n\}, \{\beta_n\}, \mathbf{V}_{\tilde{\beta}}^q \sim \mathcal{N}(\bar{\boldsymbol{\kappa}}_{\tilde{\beta}}^q, \bar{\mathbf{K}}_{\tilde{\beta}}^q) \quad (\mathbf{V}_{\tilde{\beta}}^q)^{-1}|\{e_n\}, \{\beta_n\}, \boldsymbol{\mu}_{\tilde{\beta}}^q \sim \mathcal{W}(\bar{\boldsymbol{\zeta}}_{\tilde{\beta}}^q, (\bar{\mathbf{Z}}_{\tilde{\beta}}^q)^{-1}),$$

where the posterior hyperparameters are

$$\begin{aligned} \bar{\mathbf{K}}_{\tilde{\beta}}^q &= \left( \mathbf{K}_{\tilde{\beta}}^{-1} + (\mathbf{V}_{\tilde{\beta}}^q)^{-1} \left( \sum_{n=1}^N e_n^q \right) \right)^{-1} & \bar{\mathbf{Z}}_{\tilde{\beta}}^q &= \underline{\mathbf{Z}}_{\tilde{\beta}} + \sum_{n=1}^N e_n^q (\tilde{\beta}_n - \boldsymbol{\mu}_{\tilde{\beta}}^q) (\tilde{\beta}_n - \boldsymbol{\mu}_{\tilde{\beta}}^q)' \\ \bar{\boldsymbol{\kappa}}_{\tilde{\beta}}^q &= \bar{\mathbf{K}}_{\tilde{\beta}}^q \left( \mathbf{K}_{\tilde{\beta}}^{-1} \boldsymbol{\kappa}_{\tilde{\beta}} + (\mathbf{V}_{\tilde{\beta}}^q)^{-1} \sum_{n=1}^N (e_n^q \tilde{\beta}_n) \right) & \bar{\boldsymbol{\zeta}}_{\tilde{\beta}}^q &= \underline{\boldsymbol{\zeta}}_{\tilde{\beta}} + \sum_{n=1}^N e_n^q \end{aligned}$$

with  $\boldsymbol{\kappa}_{\tilde{\beta}} = \mathbf{0}$ ,  $\mathbf{K}_{\tilde{\beta}} = 100\mathbf{I}$ ,  $\boldsymbol{\zeta}_{\tilde{\beta}} = |\tilde{\beta}|$ , and  $\underline{\mathbf{Z}}_{\tilde{\beta}} = \mathbf{I}$  the prior hyperparameters.

- The conditional posterior  $p(\beta_n, e_n|\mathbf{x}_n, \mathbf{y}_n, \{\boldsymbol{\mu}_{\tilde{\beta}}^q\}, \{\mathbf{V}_{\tilde{\beta}}^q\}, \mathbf{h}) \propto p(\mathbf{x}_n, \mathbf{y}_n|\beta_n) \cdot p(\beta_n|e_n, \{\boldsymbol{\mu}_{\tilde{\beta}}^q\}, \{\mathbf{V}_{\tilde{\beta}}^q\}) \cdot p(e_n|\mathbf{h})$  for  $(\beta_n, e_n)$  does not have a convenient form, so we draw from it using the Metropolis-Hastings algorithm. We generate a candidate  $e'_n$  using a random walk that allocates the individual to a utility type, and then we generate a candidate  $\beta'_n$ , taking care that it satisfies the constraints given  $e'_n$ . We accept/reject candidates according to the acceptance probability  $\min \left\{ \frac{p(\beta'_n, e'_n|\mathbf{x}_n, \mathbf{y}_n, \{\boldsymbol{\mu}_{\tilde{\beta}}^q\}, \{\mathbf{V}_{\tilde{\beta}}^q\}, \mathbf{h}) \cdot q(\beta_n, e_n)}{p(\beta_n, e_n|\mathbf{x}_n, \mathbf{y}_n, \{\boldsymbol{\mu}_{\tilde{\beta}}^q\}, \{\mathbf{V}_{\tilde{\beta}}^q\}, \mathbf{h}) \cdot q(\beta'_n, e'_n)}, 1 \right\}$ .

We make 5 million draws, discarding 10% as burn-in and retaining every 50<sup>th</sup> to mitigate serial correlation. Convergence is monitored by inspecting trace plots and posterior density plots. In Section IA.3 of the Internet Appendix, we present the trace plots of the population model parameters, which indicate no problems with convergence.