ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

FROM AGGREGATE BETTING DATA TO INDIVIDUAL RISK PREFERENCES

PIERRE-ANDRÉ CHIAPPORI
Department of Economics, Columbia University

BERNARD SALANIÉ
Department of Economics, Columbia University

FRANÇOIS SALANIÉ
Toulouse School of Economics, University of Toulouse Capitole and INRA

AMIT GANDHI
Department of Economics, University of Pennsylvania

FROM AGGREGATE BETTING DATA TO INDIVIDUAL RISK PREFERENCES

PIERRE-ANDRÉ CHIAPPORI
Department of Economics, Columbia University

BERNARD SALANIÉ
Department of Economics, Columbia University

FRANÇOIS SALANIÉ
Toulouse School of Economics, University of Toulouse Capitole and INRA

AMIT GANDHI
Department of Economics, University of Pennsylvania

We show that even in the absence of data on individual decisions, the distribution of individual attitudes towards risk can be identified from the aggregate conditions that characterize equilibrium on markets for risky assets. Taking parimutuel horse races as a textbook model of contingent markets, we allow for heterogeneous bettors with very general risk preferences, including non-expected utility. Under a standard single-crossing condition on preferences, we identify the distribution of preferences among the population of bettors and we derive testable implications. We estimate the model on data from U.S. races. Specifications based on expected utility fit the data very poorly. Our results stress the crucial importance of nonlinear probability weighting. They also suggest that several dimensions of heterogeneity may be at work.

KEYWORDS: Identification, revealed preferences, attitudes towards risk.

INTRODUCTION

THE LITERATURE DEVOTED TO THE EMPIRICAL ESTIMATION of individual attitudes to risk is by now quite large. To quote but a few recent examples:[1] Barsky, Juster, Kimball, and Shapiro (1997) used survey questions and observations of actual behavior to measure relative risk aversion. Results indicate that this parameter varies between 2 (for the first decile) and 25 (for the last decile), and that this heterogeneity is poorly explained by demographic variables. Guiso and Paiella (2006) reported similar findings, and used the term "massive unexplained heterogeneity." Chiappori and Paiella (2011) observed the financial choices of a sample of households across time, and used these panel data to show that while a model with constant relative risk aversion well explains each household's

Pierre-André Chiappori: pc2167@columbia.edu
Bernard Salanié: bsalanie@columbia.edu
François Salanié: francois.salanie@inra.fr
Amit Gandhi: akgandhi@sas.upenn.edu

[1]Among earlier attempts to measure the dispersion of risk aversion, one may mention for instance Binswanger (1980), Hey and Orme (1994), and the references in Wik, Kebede, Bergland, and Holden (2004). We thank an anonymous referee for suggesting some of these references.

choices, the corresponding coefficient is highly variable across households (its mean is 4.2, for a median of 1.7). Distributions of risk aversions have also been estimated using data on television games (Beetsma and Schotman (2001)), insurance markets (Cohen and Einav (2007), Barseghyan, Molinari, O'Donoghue, and Teitelbaum (2013), Barseghyan, Molinari, and Teitelbaum (2016)), or risk sharing within closed communities (Chiappori, Samphantharak, Schulhofer-Wohl, and Townsend (2014)).

These papers, and many others, rely on data on individual behavior. Indeed, a widely shared view posits that microdata are indispensable to analyze attitudes towards risk, particularly in the presence of observed or unobserved heterogeneity. The present paper challenges this claim. It argues that, in many contexts, the distribution of risk attitudes can be nonparametrically identified, even in the absence of data on *individual* decisions. We only need to use the aggregate conditions that characterize an equilibrium, provided that such equilibria can be observed on a large set of different menus of choices for the same population. The crux of our argument is that the equilibrium mapping reveals information about the distribution of risk attitudes within the population under consideration. While a related approach has often been used in other fields (e.g., empirical industrial organization), it is much less common for the estimation of a distribution of individual attitudes towards risk.[2]

In practice, we focus on "win bets" placed in horse races that use parimutuel betting. Bettors choose which horse to bet on, and those who bet on the winning horse share the total amount wagered in the race (minus the organizer's take). This has several attractive properties for our purposes. First, a win bet is simply a state-contingent asset. Second, observing the odds of a horse—the rate of return if it wins—is equivalent to observing its market share. Third, large samples of races are readily available. Finally, the decision we model is discrete (which horse to bet on), and the stochastic process that generates a win is very simple.

Each race can be represented as a menu of choices, which consists of probabilities and odds; we can simultaneously observe or at least estimate both the odds and the winning probability of each horse. Bettors choose between high return/low probability horses (longshots) and low return/high probability horses (favorites). If bettors were risk-neutral, equilibrium odds would be directly proportional to winning probabilities. In general, the mapping from probabilities to odds is more complex; it reflects the response of bettors to given menus of risky choices. If we observe a large enough number of races with enough variation in odds and winning probabilities, and the population of bettors has the same distribution of preferences in all races, then we can learn about this distribution by observing the mapping from race odds to probabilities.

We analyze two sets of questions. The first one is *testability*: given any representation of individual decision under uncertainty, does our general model generate testable restrictions on equilibrium patterns—as summarized by the relationship between probabilities and odds? And can more specific formulations (e.g., expected utility) be tested against the general model? The second is *identifiability*: under which conditions is it possible to recover the distribution of individual preferences from equilibrium patterns? Our goal here is to minimize the restrictions we a priori impose on the distribution of preferences.

We show that only four surprisingly mild assumptions are needed. The first one is that when choosing between bets, agents only consider their direct outcomes: the utility derived from betting on a horse with a given winning probability does not depend on the

---

[2]Chabi-Yo, Leisen, and Renault (2014) is a notable exception. They extended the standard CAPM analysis to allow for skewness risk. By applying small noise expansions to get closed-form formulas, the authors showed how the cross-sectional distribution of preferences maps into equilibrium prices.

characteristics of other horses. While this assumption does rule out a few existing frameworks (e.g., those based on regret theory), it is compatible with the vast majority of models of decision-making under uncertainty. Second, we assume that each agent bets the same amount in every race. Whether *nonparametric* tests and identification like those we develop below could be constructed with endogenous bet amounts or would require information at the individual level is an open question. Third, agents' decisions regarding bets are, in our model, based on the *true* distribution of winning probabilities. Note, however, that we do not impose that valuations be linear in probabilities; on the contrary, we allow for the type of probability weighting emphasized by modern decision theory, starting with Yaari's dual model or Kahneman and Tversky's cumulative prospect theory. Finally, we assume that heterogeneity of preferences is one-dimensional, and satisfies a standard single-crossing condition. The corresponding heterogeneity may affect utility, probability weighting, or both; in that sense, our framework is compatible with a wide range of theories. Our methods can also be extended to at least some forms of multidimensional heterogeneity; we address this briefly, but we leave a more general treatment for future research.

Our main theoretical result states that, under these conditions, an equilibrium always exists and is unique. We then show that we can both identify and test the model. We derive strong testable restrictions on equilibrium patterns. When these restrictions are fulfilled, we can identify the distribution of preferences in the population of bettors; in particular, we can compare various classes of preferences and distributions.

We then provide an empirical application of these results. In our setting, the concept of *normalized fear of ruin* (NF) provides the most adequate representation of the risk/return trade-off. Normalized fear of ruin directly generalizes the fear-of-ruin index introduced in an expected utility setting by Aumann and Kurz (1977). Bettors value returns (odds) as well as the probability of winning. The NF simply measures the elasticity of required return with respect to probability along an indifference curve in this space. As such, it can be defined under expected utility maximization, in which case it does not depend on probabilities; but also in more general frameworks, with probability weightings or various non-separabilities. We show that the identification problem boils down to recovering the NF index as a function of odds, probabilities and a one-dimensional heterogeneity parameter. We provide a set of necessary and sufficient conditions for a given such function to be rationalizable as an NF. These conditions provide the testable restrictions mentioned above, both for the general model and for specific versions. We also show that under these conditions, the distribution of NF is nonparametrically identified.[3]

Finally, we estimate our model on a sample of more than 25,000 races involving some 200,000 horses. Since the populations in the various "markets" must, in our approach, have similar distributions of preferences, we focus on races taking place during weekdays, on urban racetracks. Since we observe market shares, the single-crossing assumption allows us to characterize the one-dimensional index of each marginal bettor (i.e., the rank of the bettor indifferent between two horses). We specify a very general value function that depends on the winning probability, the corresponding return, and this index, based on orthogonal polynomials. We use the indifference conditions to estimate the winning probabilities and parameters by a simple log-likelihood maximization. The advantage of such a strategy is that it allows for nonparametric estimation of both a general model

---

[3]Appendix B of the Supplemental Material (Chiappori, Salanié, Salanié, and Gandhi (2019)) extends these results to the case when bettors decide which races they will bet on. For reasons discussed below, we do not explicitly consider that decision in our empirical exercise.

(involving unrestricted non-expected utility with general one-dimensional heterogeneity) and several nested submodels (including homogeneous and heterogeneous versions of expected utility maximization, Yaari's dual model, and rank-dependent expected utility).

Our empirical conclusions are quite striking. First, the type of preferences that are routinely used in the applied literature (e.g., constant relative or absolute risk aversion) are incompatible with the data. They imply restrictive conditions on the shape of the NF functions that our estimates clearly reject.[4] This suggests that the parametric approaches adopted in much applied work should be handled with care, as they may imply unduly restrictive assumptions.

Second, models relying on an expected utility (EU) framework do not perform well; their fit is quite poor, even for heterogeneous versions of the model. Moreover, single-crossing restrictions are violated for approximately half of our sample, thus casting a doubt on whether a one-dimensional index is enough to capture the impact of heterogeneity. In fact, our preferred models are relatively parsimonious versions of homogeneous rank-dependent expected utility (RDEU) preferences, and of homogeneous NEU preferences. In both cases, the main role is played by distortions of probabilities.[5] Introducing heterogeneity within the NEU framework further improves the fit, but only slightly; this last conclusion must be taken with a pinch of salt, since we may be reaching the limits of what our data can robustly say.

Related Literature

The notion that testable restrictions may be generated regarding the form of the equilibrium manifold is not new, and can be traced back to Brown and Matzkin (1996) and Chiappori, Ekeland, Kubler, and Polemarchakis (2002, 2004); the latter, in addition, introduced the idea of recovering individual preferences from the structure of the manifold. But to the best of our knowledge, these papers have not led to empirical applications. Our contributions here are most closely related to the literature on estimating and evaluating theories of individual risk preferences, and also to the literature on identification of random utility models. There is now a large literature that tests and measures theories of individual risk preference using laboratory methods (see, e.g., Camerer and Kunreuther (1989), Harless and Camerer (1994); and Bruhin, Fehr-Duda, and Epper (2010)). There is also a sizable literature that directly elicits individual risk preferences through survey questions (see, e.g., Barsky et al. (1997), Bonin, Dohmen, Falk, Huffman, and Sunde (2007), Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2011)) and correlates these measures with other economic behaviors. The literature that studies risk preferences as revealed by market transactions is much more limited. Most of it has focused on insurance choices (see, e.g., Cohen and Einav (2007), Sydnor (2010); and Barseghyan et al. (2013) and Barseghyan, Molinari, and Teitelbaum (2016)) and gambling behavior (see, e.g., Andrikogiannopoulou and Papakonstantinou (2016)). However, all of these studies fundamentally exploit individual-level demand data to estimate risk preferences and document heterogeneity.

The literature on estimating risk preferences from market-level data has almost exclusively used a representative agent paradigm. Starting with Weitzman (1965), betting

---

[4]For instance, under such commonly used representations as CARA or CRRA preferences, any given individual is either always risk-averse or always risk-loving. However, under our preferred specification, a given bettor may be risk-averse for some bets and risk-loving for others.

[5]See Barberis (2013) for a recent analysis of probability weighting,

markets have served as a natural source of data for representative agent studies of risk preferences due to the textbook nature of the gambles that are offered. In the context of racetrack betting, Jullien and Salanié (2000) and Snowberg and Wolfers (2010) provided evidence showing that a representative agent with nonlinear probability weighting better explains the pattern of prices at the racetrack as compared to an expected utility maximizing representative agent. Aruoba and Kearney (2011) presented similar findings using cross-sectional prices and quantities from state lotteries. These representative agent studies of betting markets stand in contrast to a strand of research that has emphasized belief heterogeneity as an important determinant of equilibrium in security markets. Ottaviani and Sorensen (2010) and Gandhi and Serrano-Padial (2015) argued that heterogeneity of beliefs and/or information of risk-neutral agents can explain the well-known favorite-longshot bias that characterizes many betting markets. Gandhi and Serrano-Padial furthermore estimated the degree of belief heterogeneity revealed by equilibrium patterns. In contrast, our aim here is to fully explore the consequences of heterogeneity in preferences. Specifically, we nonparametrically identify and estimate heterogeneous risk preferences from market-level data. Furthermore, our theoretical framework, while excluding heterogeneity in beliefs, allows for heterogeneity in probability weighting across agents; and our nonparametric approach allows us to compare this and other theories (such as heterogeneity in risk preferences in an expected utility framework).

Finally, our paper makes a contribution to the identification of random utility models of demand. Random utility models have become a popular way to model market demand for differentiated products following Bresnahan (1987), Berry (1994), and Berry, Levinsohn, and Pakes (1995). A lingering question in this literature is whether preference heterogeneity can indeed be identified from market-level observations alone. Along with Chiappori, Gandhi, Salanié, and Salanié (2009), our paper shows that a nonparametric model of vertically differentiated demand can be identified from standard variation in the set of products available across markets. In particular, we exploit a one-dimensional source of preference heterogeneity that satisfies a standard single-crossing condition consistent with vertically differentiated demand. We show that the identification of inverse demand from the data allows us to nonparametrically recover this class of preferences. This stands in contrast to the work by Berry and Haile (2014, 2016), which relies on a combination of index restrictions and instrumental variables.[6] We instead show identification of random utility by imposing the single-crossing structure.

We present the institution, assumptions, and the structure of market equilibrium in Section 1. In Section 2, we explain the testable restrictions on observed demand behavior implied by the model, and we show that these restrictions are sufficient to identify preferences. Section 3 describes the data, while Section 4 discusses the estimation strategy. We describe our results in Section 5, and we end with some concluding remarks. Some proofs are in the Appendix. The text also refers to the Supplemental Material (Chiappori et al. (2019)) for additional elements.

1. THEORETICAL FRAMEWORK

Parimutuel

We start with the institutional organization of parimutuel betting. Consider a race with horses $i = 1, \dots, n$. We focus on "win bets," that is, bets on the winning horse: each dollar

---

bet on horse i pays a net return of R_i dollars if horse i wins, and is lost otherwise. R_i is called the *odds* of horse i, and in parimutuel races it is determined by the following rule: all money wagered by bettors constitutes a pool that is redistributed to those who bet on the winning horse, apart from a share t corresponding to taxes and a "house take." Accordingly, if s_i is the share of the pool corresponding to the sums wagered on horse i, the payment to a winning bet of \$1 is

$$R_i + 1 = \frac{1-t}{s_i}. \tag{1}$$

Hence, odds are not set by bookmakers; instead, they are determined by the distribution $(s_1, \ldots, s_n)$ of bets among horses. Odds are mechanically low for those horses on which many bettors laid money (favorites), and they are high for longshots.[7] Since market shares sum to 1, these equations together imply

$$\frac{1}{1-t} = \sum_i \frac{1}{R_i + 1}. \tag{2}$$

Hence, knowing the odds $(R_1, \ldots, R_n)$ allows to compute both the take t and the shares in the pool $(s_1, \ldots, s_n)$.

Probabilities

We now define an *n*-horse *race* $(\mathbf{p}, t)$ by a vector of positive probabilities $\mathbf{p} = (p_1, \ldots, p_n)$ in the *n*-dimensional simplex, and a take $t \in (0, 1)$. Note that p_i is the *objective* probability that horse i wins the race. Our setting is thus compatible with traditional models of decision under uncertainty, in which all agents agree on the probabilities of the various states of the world, and these probabilities are correct. This framework singles out preferences as the driving determinant of odds; it accords well with empirical work that shows how odds reflect most relevant information about winning probabilities.[8] It is also consistent with the familiar rational expectations hypothesis; in fact, we will show that a rational expectations equilibrium exists and is unique in our setting. It is important to stress, however, that our framework is also compatible with more general models of decision-making. In particular, it allows for the type of *probability weighting* that characterizes many non-expected utility functionals, whereby the actual decision process may involve arbitrary increasing functions of the probabilities. Moreover, these probability weights may be agent-specific, as we shall see. In other words, our general framework encompasses both "traditional" models, in which agents always refer to objective probability and heterogeneity only affects preferences, and more general versions in which different agents weigh probabilities differently. The only strong restriction we will impose bears on the *dimension* of the heterogeneity, not on its nature.

Following the literature to date,[9] we endow each bettor with a standardized bet amount that he allocates to his most preferred horse in the race. In particular, we do not allow participants to bet heterogeneous amounts. Therefore, the shares (s_i) in the pool defined above can be identified to market shares. Any bettor looks on a bet on horse i as a lottery

---

[7]According to this formula, odds can even be negative, if s_i is above $(1 - t)$; it never happens in our data.

[8]See Sung and Johnson (2008) and Ziemba (2008) for recent surveys on the informational efficiency of betting markets.

[9]See Weitzman (1965), Jullien and Salanié (2000), Snowberg and Wolfers (2010), among others.

that pays R_i with probability p_i, and pays (-1) with probability $(1 - p_i)$. We denote this lottery by (R_i, p_i), and call it a *gamble*. By convention, throughout the paper we index horses by decreasing probabilities ($p_1 > \cdots > p_n > 0$), so that horse 1 is the favorite.

Risk-Neutral Bettors

As a benchmark, consider the case when bettors are risk-neutral, and thus only consider the expected gain associated to any specific bet.[10] Equilibrium then requires expected values to be equalized across horses. Since bets (net of the take) are redistributed, this yields

$$p_i R_i - (1 - p_i) = -t,$$

which, together with (2), gives probabilities equal to

$$p_i^n(R_1, \ldots, R_n) = \frac{\dfrac{1}{R_i + 1}}{\displaystyle\sum_j \dfrac{1}{R_j + 1}} = \frac{1 - t}{1 + R_i} = s_i. \tag{3}$$

By extension, for any set of odds $(R_1, \ldots, R_n)$, we will call the above probabilities p_i^n *risk-neutral probabilities*. These probabilities are exactly equal to the shares s_i in the betting pool, as defined in (1) and (2). Many stylized facts (for instance, the celebrated *favorite-longshot bias*) can easily be represented by comparing the "true" probabilities with the risk-neutral ones—more on this below.

1.1. *Preferences Over Gambles*

We consider a continuum of bettors, indexed by a parameter θ. Each bettor θ is characterized by a valuation function $V(R, p, \theta)$, defined over the set of all possible gambles (R, p). In a given race, θ bets on the horse i that gives the highest value to $V(R_i, p_i, \theta)$. As usual, $V(\cdot, \cdot, \theta)$ is only defined ordinally, that is, up to an increasing transform. We consequently normalize to zero the value $V(-1, p, \theta) \equiv 0$ of losing 1 with certainty.

Note that each $V(\cdot, \cdot, \theta)$ is a utility function defined *on the space of gambles*. As such, it is compatible with expected utility, but also with most non-expected utility frameworks; one goal of this paper is precisely to compare the respective performances of these various models on our data. Finally, the main restriction implicit in our assumption is that the utility derived from betting on a given horse does not depend on the other horses in the race; we thus rule out models based, for instance, on regret theory,[11] and more generally, any framework in which the valuation of a bet depends not only on the characteristics of the bet but also on the whole set of bets available.

We will impose several assumptions on V. We start with very weak ones:

ASSUMPTION 1: *For each θ, $(R, p) \mapsto V(R, p, \theta)$ is continuously differentiable almost everywhere; and it is increasing with R and p.*

---

[10]Clearly, a risk-neutral player will not take a bet with a negative expected value unless she derives some fixed utility from gambling (see Conlisk (1993)). The assumption we maintain in this paper is that this "utility of gambling" does not depend on the particular horse on which the bet is placed: *conditional on betting*, bettors still select the horse that generates the highest expected gain.

[11]See, for example, Gollier and Salanié (2006).

Differentiability is not crucial; its only role is to simplify some of the equations. Our framework allows for a kink at some reference point, for instance, as implied by prospect theory. The second part of the assumption reflects first-order stochastic dominance: bettors prefer bets that are more likely to win, or that have higher returns when they do. We now introduce another technical requirement, which we will use when proving the existence of a rational expectations equilibrium:

ASSUMPTION 2—*For any θ, R, $p > 0$:*
- *for any $p' > 0$, there exists R' such that $V(R, p, \theta) < V(R', p', \theta)$;*
- *for any R', there exists $p' > 0$ such that $V(R, p, \theta) > V(R', p', \theta)$.*

Assumption 2 is very weak: it only requires that any higher return can be compensated by a lower probability of winning, and vice versa.

The Normalized Fear of Ruin (NF)

The trade-off between risk and return is crucial in decision-making under uncertainty; and we aim to quantify it using observed choices. This trade-off can be described in several ways. One is the marginal rate of substitution[12] w:

$$w(R, p, \theta) \equiv \frac{V_p}{V_R}(R, p, \theta) > 0.$$

Since each utility function $V(\cdot, \cdot, \theta)$ is only defined up to an increasing transform, the properties of w fully determine the bettors' choices among gambles. We chose to focus on a slightly different index, which we call the *normalized fear of ruin* (NF):

$$\mathrm{NF}(R, p, \theta) \equiv \frac{p}{R+1} \frac{V_p}{V_R}(R, p, \theta) = \frac{p}{R+1} w(R, p, \theta) > 0.$$

Using NF rather than more traditional measures of risk-aversion has several advantages. It is unit-free, as it is the elasticity of required return with respect to probability on an indifference curve:

$$\mathrm{NF} = -\left.\frac{\partial \log(R+1)}{\partial \log p}\right|_V.$$

As such, it measures the local trade-off between risk and return. Moreover, it has a "global" interpretation for the type of binomial lotteries we are dealing with. The NF index of a risk-neutral agent is identically equal to 1. An index above 1 indicates that the agent is willing to accept a lower expected return $p(R + 1) - 1$ in exchange for an increase in the probability p. Conversely, if an agent with an index below 1 is indifferent between betting on a favorite (p, R) and a longshot $(p' < p, R' > R)$, then it must be that the expected return on the longshot is below that on the favorite. For instance, in a representative agent context, the favorite-longshot bias can be explained by the representative agent having a NF index below 1. However, our approach allows for heterogeneous bettors and can accommodate the existence of bettors with different NF indices.

---

[12]Throughout the paper, we use subscripts to indicate partial derivatives.

FIGURE 1.—The fear of ruin.

The Expected Utility Case

In an expected utility framework, the NF index has a simple expression. With a zero utility $u(-1, \theta) = 0$ from losing the bet, we have that

$$V(R, p, \theta) = pu(R, \theta),$$

and therefore

$$\text{NF}(R, p, \theta) = \frac{1}{R+1} \frac{u}{u_R}(R, \theta),$$

so that the NF index[13] is independent from the probability p. Geometrically, $\text{NF}(R)$ is the ratio of two slopes on the graph of the utility function: that of the chord linking the points $(-1, 0)$ (losing the bet) and $(R, u(R))$ (winning it), and that of the tangent to the utility graph at $(R, u(R))$ (see Figure 1).

The properties of the NF index in the expected utility case are well-known.[14] A sufficient condition for an agent to have a higher NF index than another agent at all values of R is that the former be more risk-averse than the latter.[15] Consequently, if the agent is risk-averse, then his NF index is larger than 1; if he is risk-loving, it is smaller than 1.

While the NF index need not be monotonic in R, specific functional forms generate additional properties. For example, an agent with constant absolute risk-aversion is either risk-averse (her $\text{NF}(R)$ is above 1 and increasing) or risk-loving (and then her NF is below 1 and decreasing). The same "fanning out" holds with constant relative risk-aversion. These are testable predictions; and as we shall see, our more flexible estimates show NF indices that tend to fan in rather than fan out, and sometimes cross the $\text{NF} = 1$ line.

---

[13]The ratio u/u_R was called the *fear-of-ruin* index by Aumann and Kurz (1977)—hence our choice of the name for NF.

[14]See Foncel and Treich (2005).

[15]Recall that u is more risk-averse than v if there exists an increasing and concave function k such that $u = k(v)$. Given our normalization $u(0) = v(0) = 0$, this implies that k is such that $k(x)/x$ decreases with x. This property is equivalent to u having a higher NF index than v at any value of R (Foncel and Treich (2005)).

1.2. *Single-Crossing Assumption*

Our next assumption imposes a single-crossing property that drives our approach to identification.

ASSUMPTION 3—Single-crossing: 1. *The heterogeneity parameter θ is a scalar.*
2. *Consider two gambles (R, p) and (R', p'), with $p' < p$. If, for some θ, we have*

$$V(R, p, \theta) \leq V(R', p', \theta),$$

then for all $\theta' > \theta$,

$$V(R, p, \theta') < V(R', p', \theta').$$

Given first-order stochastic dominance as per Assumption 1, if θ prefers the gamble with the lowest winning probability ($p' < p$), then it must be that its odds are higher ($R' > R$), so that the gamble (R', p') is riskier. Assumption 3 states that if θ prefers the riskier gamble, any agent θ' above θ will, too. The single-crossing assumption thus imposes that agents can be sorted according to their "taste for risk": higher θ's prefer longshots, while lower θ's prefer favorites.

Assumption 3 has a well-known differential characterization, which we state without proof:[16]

LEMMA 1: *Suppose that V is differentiable everywhere on some open set $\mathcal{O}$. Then Assumption 3 holds on $\mathcal{O}$ if and only if, for any (R, p, θ) in $\mathcal{O}$, the marginal rate of substitution $w(R, p, \theta)$, or equivalently, the normalized fear-of-ruin index $\mathrm{NF}(R, p, \theta)$, is decreasing in θ.*

Since Assumption 3 only refers to an ordering of θ, without loss of generality we normalize θ to be uniformly distributed on the interval $[0, 1]$. This essentially makes θ a quantile of the distribution of "preference for riskier bets." The precise scope of the single-crossing condition can be better seen on a few examples.

1. Expected Utility

As above, we normalize to zero the utility of losing the bet, so that

$$V(R, p, \theta) = pu(R, \theta).$$

Single-crossing holds if and only if the normalized fear of ruin is decreasing in θ. A sufficient condition is that lower θ's be more risk-averse at any value of R. For instance, in the CARA case, consider a population of bettors indexed by their absolute risk-aversion λ:

$$u(R, \lambda) = \frac{\exp(\lambda) - \exp(-\lambda R)}{\lambda},$$

where λ has a c.d.f. F_λ. Then

$$\mathrm{NF}(R, p, \lambda) = \frac{1}{R + 1} \frac{u}{u_R} = \frac{\exp(\lambda(1 + R)) - 1}{\lambda(1 + R)}$$

---

[16]See, for instance, Athey (2001, 2002).

increases with λ. If we define $\theta = 1 - F_\lambda(\lambda)$, then, by construction, θ is uniformly distributed on $[0, 1]$ and NF decreases in θ, so that Assumption 3 holds. Note also that NF is in fact an increasing function of $\lambda(R + 1)$, so that the normalized fear of ruin "fans out": for any given λ (or θ), it moves away from the risk-neutral level of 1 as odds increase.

A similar result holds for CRRA functions, with an additional twist. If $W > 1$ denotes the agent's wealth, then easy calculations give

$$\mathrm{NF}(R, p, \gamma, W) = \frac{W + R}{1 - \gamma}\left(1 - \left(\frac{W - 1}{W + R}\right)^{1 - \gamma}\right),$$

which is increasing in the relative risk-aversion index γ. Again, Assumption 3 holds if we define θ as $1 - F_\gamma(\gamma)$ where F_γ is the c.d.f. of γ. However, an agent's choice now also depends on the agent's wealth. The latter can be seen as either an alternative or an additional source of heterogeneity. That is, we can model a population of bettors with identical relative risk-aversion γ but different initial wealth W and define $\theta = F_W(W)$; since $\mathrm{NF}(R, p, \gamma, W)$ is also increasing in W, the model will still satisfy the single-crossing assumption. More ambitiously, we could move beyond Assumption 3 and consider agents who differ in both wealth and risk-aversion, generating bi-dimensional heterogeneity; we will return to this issue in Section 2.3.

This point is more general: in the absence of individual data, we cannot possibly distinguish between heterogeneity in "preferences," in wealth, and in background risk, for instance; we can only estimate the resulting heterogeneity in attitudes towards lotteries.

2. Rank-Dependent Expected Utility Theory

RDEU enriches the previous framework by allowing for a nonlinear weighting of probabilities: the utility V can be written

$$V(R, p, \theta) = G(p, \theta)u(R, \theta).$$

For Assumption 1 to hold, the probability weighting function G must increase in p and the utility function u must increase in R. In general, both functions may vary with θ.

Now remember that V is only defined up to an increasing transformation: we can only hope to identify its indifference curves, whose slope is NF. In the RDEU case, the NF index is a product of two terms:

$$\mathrm{NF}(R, p, \theta) = \frac{p}{R + 1}\frac{V_p}{V_R} = \frac{1}{R + 1}\frac{u}{u_R}\frac{pG_p}{G}. \tag{4}$$

The first term is the NF index for an expected utility maximizer with utility u, which is the elasticity of $u(R, \theta)$ with respect to the gross return $(R + 1)$. The second term is the elasticity of $G(p, \theta)$ with respect to the probability p of a win. It is the NF index that would obtain if u were linear in R for all θ, as in the "dual expected utility" model of Yaari (1987):

$$V(R, p, \theta) = G(p, \theta)(R + 1).$$

For Yaari-like preferences, the NF index is independent of R; and single-crossing requires that G_p/G, which is positive, be decreasing in θ. In words, this means that larger θ's put more weight on small probabilities. Again, this allows us to account for some heterogeneity in beliefs. Note that since Yaari's model sets the elasticity of u to 1, we can identify the

elasticity of G and its variations with θ in this more restricted model. However, with only one dimension of variation in θ, it is difficult to account for heterogeneity in both u and G simultaneously in the RDEU model. What we can and will do is test heterogeneity in u versus heterogeneity in G.

3. Extensions and Limitations

Many other families of preferences, such as cumulative prospect theory, also fit within our setting—although the single-crossing condition becomes more complicated. Others may only be accommodated under some restrictions. For instance, the reference-dependent theory of choice under risk of Köszegi and Rabin (2007) yields a choice functional that fits our framework as long as it respects stochastic dominance. We could also incorporate ambiguity-aversion in the "exponential tilting" form introduced by Hansen and Sargent (e.g., in their 2007 book, or Hansen (2007)). However, in our very simple choice problems with static decision-making, it is observationally equivalent to increased risk-aversion.[17]

Our approach has two main limitations. First, we require that agents only pay attention to realized consequences. Some models of decision under uncertainty relax this assumption; regret theory and disappointment-aversion, for instance, are only compatible with our setting in restricted cases. Second, we only allow for one dimension of heterogeneity. Our approach is only compatible with models involving heterogeneity in both preferences and beliefs if these two dimensions are governed by the same parameter. Multidimensional nonparametric heterogeneity is a very difficult problem, which is left for future work; for the time being, we shall simply provide a short discussion of a possible approach (see Section 2.3).

1.3. Market Shares and Equilibrium

The winning probabilities $\mathbf{p}$ and the take t are assumed exogenous and characterize a race. In contrast, the odds are endogenous: the bettors' behavior determines market shares, which in turn determine odds through the parimutuel rule (1). In this setting, it is natural to rely on the concept of rational expectations equilibria: agents determine their behavior given their anticipations on odds, and these anticipations are fulfilled in equilibrium. We now show that for our framework, a rational expectations equilibrium exists. Moreover, our characterization of the equilibrium condition in terms of the single-crossing assumption will provide the key to the identification of preferences.

Focus on a given race with n horses, and assume that the win probabilities $\mathbf{p}$ and odds $\mathbf{R}$ are given, and known to all agents. Each agent then optimizes on which horse to bet on. As a simple consequence of the single-crossing condition, the choices bettors make partition them into a sequence of intervals:

LEMMA 2: *Suppose that $\mathbf{p}$ and $\mathbf{R}$ are such that all market shares are positive: $s_i > 0$ for all $i = 1, \ldots, n$. Then, under Assumptions 1, 2, and 3, there exists a family $(\theta_j)_{j=0,\ldots,n}$, with $\theta_0 = 0 < \theta_1 < \cdots < \theta_{n-1} < \theta_n = 1$, such that:*

• *for all $i = 1, \ldots, n$, if $\theta_{i-1} < \theta < \theta_i$, then bettor θ strictly prefers to bet on horse i than on any other horse;*

---

[17]If we let ambiguity-aversion depend on observables (for instance, it may be more prominent in races with younger horses), then we could distinguish it from risk-aversion. Gandhi and Serrano-Padial (2015) made use of a related strategy.

- *for all $i = 1, \ldots, n-1$, we have*

$$V(p_i, R_i, \theta_i) = V(p_{i+1}, R_{i+1}, \theta_i). \tag{5}$$

Lemma 2 states that if we rank horses by increasing odds in a race, bettors will self-select into n intervals; in each interval, all bettors bet for the same horse. The bounds of the intervals are defined by an indifference condition: for $i = 1, \ldots, n-1$, there exists a *marginal bettor* θ_i who is indifferent between betting on horses i and $i+1$. As a simple corollary and since we normalized the distribution of θ to be uniform, the market share s_i of horse $i = 1, \ldots, n$ is

$$s_i = \theta_i - \theta_{i-1},$$

which yields

$$\theta_i = \sum_{j=1}^{i} s_j.$$

Recall that odds are determined from market shares as in (1) and (2); therefore, in equilibrium, one must have $\theta_i = \theta_i(\mathbf{R})$, where

$$\theta_i(\mathbf{R}) \equiv \frac{\sum_{j \leq i} \dfrac{1}{R_j + 1}}{\sum_j \dfrac{1}{R_j + 1}}, \quad i = 1, \ldots, n. \tag{6}$$

At a rational expectations market equilibrium, bettors must choose optimally given odds and probabilities, as expressed in (5); and odds must result from market shares, which is what the equalities $\theta_i = \theta_i(\mathbf{R})$ impose. This motivates the following definition:

DEFINITION 1: Consider a race $(\mathbf{p}, t)$. $\mathbf{R} = (R_1, \ldots, R_n)$ is a family of equilibrium odds if and only if (2) holds and

$$\forall i < n \quad V\big(p_i, R_i, \theta_i(\mathbf{R})\big) = V\big(p_{i+1}, R_{i+1}, \theta_i(\mathbf{R})\big). \tag{7}$$

We then prove existence and uniqueness. The result is a particular instance of a more general result in Gandhi (2006), but its proof in our setting is quite direct:

PROPOSITION 1: *Under Assumptions 1, 2, and 3, for any race $(\mathbf{p}, t)$, there exists a unique family $-t < R_1 \leq \cdots \leq R_n$ of equilibrium odds.*

Hence, to each race, one can associate a unique rational expectations equilibrium, with positive market shares. This result gives a foundation to our assumption that bettors share common, correct beliefs. From an empirical viewpoint, however, only the odds are directly observable; probabilities have to be estimated. Fortunately, probabilities can be uniquely recovered from odds:

PROPOSITION 2: *Under Assumptions 1, 2, and 3, for any $\mathbf{R}$ ranked in increasing odds, there exists a unique race $(\mathbf{p}, t)$ such that $\mathbf{R}$ is a family of equilibrium odds for $(\mathbf{p}, t)$.*

As already observed, the rules of parimutuel betting allow us to infer the value of the track take t from odds, using (2). On the other hand, the relationship between odds and probabilities results from preferences. The function $\mathbf{p}(\mathbf{R}) = (p_1(\mathbf{R}), \ldots, p_n(\mathbf{R}))$ implicitly defined in Proposition 2 thus conveys some information on the underlying preferences of bettors. Since choices are fully determined by the marginal rates of substitution $w = V_{\mathbf{R}}/V_{\mathbf{p}}$, we shall say hereafter that $\mathbf{p}(\mathbf{R})$ characterizes any market equilibrium associated to V. Finally, Propositions 1 and 2 extend in a straightforward manner to the homogeneous case in which bettors are identical: each bettor must then be indifferent among all horses, in the spirit of Jullien and Salanié (2000).

2. TESTABLE IMPLICATIONS AND IDENTIFIABILITY

Assume for the moment that we observe the same population of bettors[18] faced with a large number of races $(\mathbf{p}, t)$. In each race, individual betting behavior leads to equilibrium odds $\mathbf{R}$ and market shares s, which are observable; we also observe the identity of the winning horse for each race. We also assume that the relationship between winning probabilities and equilibrium odds $\mathbf{p}(\mathbf{R})$ is known. In fact, this relationship could be estimated very flexibly from a rich enough data set; but our actual estimation strategy, which we expose in Section 4, will not rely on such a direct estimation of probabilities.

We focus here on the empirical content of our general framework. Specifically, we consider two questions. One is testability: does the theory impose testable restrictions on the form of the function $\mathbf{p}(\mathbf{R})$? The second issue relates to identifiability: given $\mathbf{p}(\mathbf{R})$, is it possible to uniquely recover the distribution of individual preferences, that is, in our setting, the normalized fear of ruin $\mathrm{NF}(R, p, \theta)$? We shall now see that the answer to both questions is positive.

2.1. *Testable Implications*

We start with testability. Since V increases in p, we can define Γ as the inverse of V with respect to p:

$$\forall R, p, \theta \quad \Gamma\big(V(R, p, \theta), R, \theta\big) = p.$$

One can then define a function G as

$$G\big(R, p, R', \theta\big) = \Gamma\big(V(R, p, \theta), R', \theta\big). \tag{8}$$

In words, $G(R, p, R', \theta)$ is the winning probability p' that would make a gamble (R', p') equivalent, for bettor θ, to the gamble (R, p). Now we can rewrite the equilibrium conditions in Definition 1 as

$$\forall i < n \quad p_{i+1}(\mathbf{R}) = G\big(R_i, p_i(\mathbf{R}), R_{i+1}, \theta_i(\mathbf{R})\big), \tag{9}$$

where $\theta_i(\mathbf{R})$ was defined in (6). We immediately obtain several properties of G:

PROPOSITION 3: *Let Assumptions 1, 2, and 3 hold. If* $\mathbf{p}(\mathbf{R})$ *characterizes market equilibria associated to some family* V, *then there exists a function* $G(R, p, R', \theta)$ *such that*

---

[18]Strictly speaking, we only need the distribution of preferences to be constant. Appendix B of the Supplemental Material considers the case when bettors endogenously decide whether to bet in each race, so that the distribution of preferences of bettors in a race depends on some of its characteristics. We provide assumptions under which the identification result still holds.

(i) *G is continuously differentiable, increasing with R and p, decreasing with θ if $R' > R$, and decreasing with R';*
 (ii) G_p/G_R *is independent of R';*
(iii) $G(R, p, R, \theta) = p$;
(iv) (9) *holds for any family $R_1 < R_2 < \cdots < R_n$.*

Of the four properties in Proposition 3, (ii) and (iv) are the main restrictions that our theory imposes on observed odds and probabilities. Property (iv) states that the winning probability $p_{i+1}(\mathbf{R})$, which could depend on the *whole family* of odds $(R_1, \ldots, R_n)$, can be computed from only four numbers: the pair of odds R_i and R_{i+1}, the index of the marginal consumer $\theta_i(\mathbf{R})$ (which can be directly inferred from market shares, as argued above), and the probability $p_i(\mathbf{R})$ of the horse ranked by bettors just above $(i + 1)$. Hence, $p_i(\mathbf{R})$ and $\theta_i(\mathbf{R})$ are sufficient statistics for the $(n-2)$ odds that are missing from this list. Moreover, G does not depend on the index i, on the number of horses n, nor on the take t. Finally, property (ii) dramatically restricts the variation in G. These and the other two properties of G listed in Proposition 3 will provide directly testable predictions of our model.[19]

2.2. *Exhaustiveness and Identification*

Take some function $\mathbf{p}(\mathbf{R})$ that satisfies the conditions we just derived. Pick a particular i (say, $i = 1$); then, by Proposition 3, for each race and each horse i, $p_{i+1}(\mathbf{R})$ can only depend on the four variables $(R_i, p_i(\mathbf{R}), R_{i+1}, \theta_i(\mathbf{R}))$. The corresponding relationship nonparametrically identifies the function G and generates a first set of testable restrictions; a second set follows from the fact that the resulting G does not depend on the choice of i.

We now show that the four properties in Proposition 3 are sufficient. From any $\mathbf{p}(\mathbf{R})$ associated to a G function that satisfies all four properties, we can recover a function $\text{NF}(R, p, \theta)$ such that $\mathbf{p}(\mathbf{R})$ characterizes the market equilibria associated to any risk preferences V whose normalized fear of ruin is NF. In turn, recovering the normalized fear of ruin NF allows to nonparametrically identify preferences, that is, to ordinally identify the function V. Specifically, the following holds:

PROPOSITION 4: *Suppose that the function $\mathbf{p}(\mathbf{R})$ satisfies the restrictions in (9) for some function G. Let S_4 be the domain over which (9) defines G, and assume that properties* (i)–(iii) *in Proposition 3 hold for G over S_4. Define S_3 to be the set of (R, p, θ) such that (R, p, R', θ) belongs to S_4 for some $R' > R$.*

Then there exists a unique (*up to increasing transforms*) *function $V(p, R, \theta)$ defined on S_3 such that $\mathbf{p}(\mathbf{R})$ characterizes the market equilibria associated to V.*

Moreover, V verifies the single-crossing property, and its normalized fear of ruin NF is

$$\text{NF}(R, p, \theta) = \frac{p}{R + 1} \frac{G_p}{G_R}(R, p, R', \theta).$$

From an empirical viewpoint, Proposition 4 proves two results. First, the properties (i)–(iv) stated in Proposition 3 are in fact sufficient: since they are strong enough to ensure the existence of a family V satisfying our assumptions, no other testable implications can

---

[19]In the homogeneous case in which bettors are identical, the results extend directly, with the only change that G does not depend on θ anymore. Similar statements are valid for the following results.

be found. Second, the MRS function w is uniquely identified. Indeed, for (8) to hold, it must be that

$$\frac{V_p}{V_R}(R, p, \theta) = \frac{G_p}{G_R}(R, p, R', \theta) \quad \text{for all } R',$$

which property (ii) of Proposition 3 makes possible. This defines w (and NF) uniquely, and consequently, the family V is identified up to an increasing function of θ. Hence, *under our assumptions, aggregate data are enough to recover heterogeneous individual preferences without any parametric assumption.*[20]

Proposition 4 qualifies this conclusion in one respect: identification only holds on the support S_3 of the random variables that we defined. This has an important consequence in our setting. Assume that no race has more than n horses. The favorite in each race, by definition, has the largest market share, and so we will always observe $\theta_1 > 1/n$. Since identification relies on boundary conditions in the θ_i's, it follows that we cannot hope to recover the family of functions $V(\cdot, \cdot, \theta)$ for $\theta < 1/n$. (More formally, the set S_3 contains no point (R, p, θ) with $\theta < 1/n$.)

2.3. *Multidimensional Heterogeneity*

We assumed so far that heterogeneity was one-dimensional, and could be described by a single parameter θ. As already mentioned in Section 1.2, we might want to go beyond this and allow for more dimensions of heterogeneity. We sketch here how it can be done in two dimensions. Assume that agents differ in two characteristics, described by scalar parameters θ and η. Without loss of generality, we normalize the marginal distributions of θ and η to be independent uniform distributions over [0, 1]. This can be done by applying the quantile transforms $\theta' = F_\theta(\theta)$ and $\eta' = F_{\eta|\theta}(\eta|\theta)$.

Now consider two gambles (R, p) and (R', p') such that $p' < p$ (therefore, $R' > R$). We impose the single-crossing property in each of the two dimensions. That is, if, for some (θ, η), we have

$$V(R, p, \theta, \eta) < V(R', p', \theta, \eta),$$

then for all $\theta' > \theta$,

$$V(R, p, \theta, \eta) < V(R', p', \theta', \eta),$$

and for all $\eta' > \eta$,

$$V(R, p, \theta, \eta) < V(R', p', \theta, \eta').$$

The interpretation is as before: for any given η, higher θ's prefer longer shots; and similarly, for any given θ, higher η's prefer longer shots. These conditions imply that the equation for the marginal bettor(s) (θ, η)

$$V(R, p, \theta, \eta) = V(R', p', \theta, \eta)$$

implicitly defines a function

$$\eta = \phi(\theta; R, p, R', p')$$

that can be represented by a decreasing curve in the (θ, η) plane. Figure 2 shows how bettors select horses in the plane for a three-horse race $R_1 < R_2 < R_3$. The main difficulty

---

[20]Appendix A of the Supplemental Material specializes this result to the case of expected utility.

FIGURE 2.—Market shares with two-dimensional heterogeneity ($n = 3$).

is that the total market share of horse $i = 2, \ldots, n - 1$ is given by the area between two isolevel curves of the function ϕ, so that for $i < n$, the cumulative market shares are

$$S_i = \sum_{j=1}^{i} s_j = \int_0^1 \phi(\theta; R_i, p_i, R_{i+1}, p_{i+1})\, d\theta$$

and we can only identify (and in fact overidentify) the function

$$\Phi(R, p, R', p') \equiv \int_0^1 \phi(\theta; R, p, R', p')\, d\theta = E(S_i | R_i = R, p_i = p, R_{i+1} = R', p_{i+1} = p').$$

This contrasts with the one-dimensional case: if θ is irrelevant, then these isolevel curves are horizontal lines and market shares *directly* identify the relevant structure of the distribution of preferences. Identification is still possible under specific assumptions regarding the form of the heterogeneity. Assume, for instance, that

$$V(R, p, \theta, \eta) = p \exp(\theta a(R, p) + \eta b(R, p)),$$

or, equivalently (since any increasing transformation can be applied to V),

$$\log V(R, p, \theta, \eta) = \log p + \theta a(R, p) + \eta b(R, p). \tag{10}$$

This quasi-linear representation, while restrictive, is standard in industrial organization and in contract theory. It is easy to see that if the functions a and b are increasing in R and do not vary too much in p (deviations from expected utility are small), then bettors with higher θ and η bet on longer shots. Moreover, the function ϕ is affine:

$$\phi(\theta; \mathbf{R}, \mathbf{p}, \mathbf{R}', \mathbf{p}') = -\frac{1}{b(\mathbf{R}', \mathbf{p}') - b(\mathbf{R}, \mathbf{p})} \left(\theta(a(\mathbf{R}', \mathbf{p}') - a(\mathbf{R}, \mathbf{p})) + \log \frac{p}{p'} \right),$$

and we can integrate over θ to get the following equation:

$$S_i = \Phi(R_i, p_i, R_{i+1}, p_{i+1}) = -\frac{\frac{1}{2}(a(R_{i+1}, p_{i+1}) - a(R_i, p_i)) + \log \dfrac{p_i}{p_{i+1}}}{b(R_{i+1}, p_{i+1}) - b(R_i, p_i)}. \tag{11}$$

This equation identifies (in fact overidentifies) the functions a and b, up to some constants that are irrelevant for the identification of preferences.[21]

3. DATA

Our entire analysis up to now has assumed a stable family of preferences $V(R, p, \theta)$ across the races in the data. This family of preferences can change with observable co-variates X, and thus we should interpret the analysis up to now as being done conditional on X.

The race data consist of a large sample of thoroughbred races (the dominant form of organized horse racing worldwide) in the United States, spanning the years 2001 through 2004. The data were collected by professional handicappers from the racing portal paceadvantage.com, and a selection of the race variables that they collect were shared with us. In particular, for each horse race in the data, we have the date of the race, the track name, the race number in the day, the number of horses in the race, the final odds for each horse, and finishing position for each horse that ran. Excluded from the data are variables that the handicappers use for their own competitive purpose, such as various measures of the racing history of each horse.

For the present analysis, we focus on the data from year 2001. For this year, we have races from 77 tracks spread over 33 states. There were 100 races in which at least one horse was "purse only," meaning that it ran but was not bet upon and hence was not assigned betting odds. In 461 races, two horses were declared winners; and in three races, there was no winner. After eliminating these three small subsamples, we had 447,166 horses in 54,169 races, an average of about 8.3 horses per race. Figure 3 shows that almost all races have 5 to 12 horses. We eliminated the other 606 races. We also dropped 44 races


FIGURE 3.—Frequency distribution of the number of horses per race.

---

[21]Appendix C of the Supplemental Material includes a proof of this claim.

TABLE I

CHARACTERISTICS OF THE SAMPLE

	Number of Horses in Race	Odds	θ
Min	5	0.05	0.145
P25	7	3.70	0.623
P50	8	8.10	0.851
P75	10	18.80	0.964
Max	12	200.00	1.000

in which one horse has odds larger than 200—a very rare occurrence. That leaves us with a sample of 442,636 horses in 53,523 races.

Table I gives some descriptive statistics. The betting odds over horses in the data range from extreme favorites (odds equaling 0.05, i.e., horses paying 5 cents on the dollar), to extreme longshots (odds equaling 200, i.e., horses paying 200 dollars on the dollar). The mean and median odds on a horse are 15.23 and 8.10, respectively: the distribution of odds is highly skewed to the right. In our sample, 18.3% of horses have $R \geq 25$ (odds of 20 or more), 6.2% of horses have $R \geq 50$, but only 0.7% have $R \geq 100$. Also, the race take (t in our notation) is heavily concentrated around 0.18: the 10th and 90th percentiles of its distribution are 0.165 and 0.209.

Figure 4 plots the raw distribution of odds up to $R = 100$. It seems fairly regular, with a mode at odds of $R = 2.5$; but this is slightly misleading. Unlike market shares, odds are


FIGURE 4.—Distribution of odds, $R \leq 100$.

TABLE II

NUMBER OF HORSES IN EACH SUBSAMPLE

	Weekday	Weekend
Rural	42,027	31,663
Urban	216,802	152,144

not a continuously distributed variable: they are rounded at the track. This rounding is of no consequence for the econometric methods we use in this paper.[22]

We built two 0–1 covariates. The first one uses the date at which a race was run to separate weekday and weekend races. To build our second covariate, we hand-collected the zip code of each racetrack, and we used it to classify each track on an urban/rural scale, thanks to the 2000 Rural-Urban Commuting Area Codes classification of the Census Bureau. Thus, our two main covariates for a race are Weekend/Weekday and Urban/Rural. Table II shows that most races are run in an urban setting, and slightly more on weekdays than on weekends. In order to focus on a relatively homogeneous sample, the results we report in the rest of this paper were obtained on the largest subsample: the 26,525 races run on weekdays in an urban setting, with 216,802 horses.

4. ESTIMATION STRATEGY

The fundamental equation of our model can be seen as determining all win probabilities recursively in any given race:

$$\forall\, i < n \quad p_{i+1}(\mathbf{R}) = G\big(R_i, p_i(\mathbf{R}), R_{i+1}, \theta_i(\mathbf{R})\big). \tag{12}$$

In this relationship, the odds $\mathbf{R}$ and interval limits $\boldsymbol{\theta}$ are directly recovered from the data. Our empirical strategy aims at estimating both the probabilities $\mathbf{p}$ and the function G.

If the value function V is known (possibly up to some parameters), then the function G can be derived from (8). Then, for each race, the system of equations (12), along with the adding up constraint $\sum_{i=1}^{n} p_i = 1$, allow us to compute the winning probabilities and therefore the likelihood of the event that the observed winner has indeed won the race. Maximizing the likelihood over all races then yields estimates of the relevant parameters. Note that this approach is a direct generalization of Jullien and Salanié (2000), in two directions: we consider heterogeneous preferences, instead of assuming homogeneous bettors; and we represent preferences in a much more flexible manner than that paper, which focused on specific classes of EU (CARA, HARA) and non-EU (RDEU, CPT) functions.

The flexible form we use for the function V is based on orthogonal polynomials. To motivate it, let us start with the benchmark of risk-neutrality. Then V coincides with the expected gain, up to normalization. We choose again to normalize the utility of losing the bet to zero, and we normalize the utility of winning to be 1 when odds are $R_M = 6$, which is close to the median odds for the marginal bettors in the sample. Then risk-neutrality would give

$$V(R, p, \theta) = p\,\frac{R+1}{R_M+1}.$$

---

[22]See Appendix D of the Supplemental Material for more information on rounding.

This suggests that a flexible generalization could use the following form:

$$V(R, p, \theta) = p\frac{R+1}{R_M+1}\exp\big(K(R, p, \theta)\big),$$

with the specification

$$K(R, p, \theta) = \sum_{k,m\geq 0} \alpha_{km}A_k(R)B_m(p) + \sum_{k,l\geq 1;m\geq 0} \beta_{klm}P_k(R_\theta)Q_l(\theta)B_m(p). \qquad (13)$$

The first sum corresponds to the homogeneous case for which all types share the same utility functions. In this formula, (A_k) is a family of polynomials chosen to be orthogonal over the distribution of R; we normalize them so that $A_k(R_M) = 0$ for all degrees $k \geq 0$. Departures from the expected utility case are captured by the family (B_m) of polynomials, which we chose to be orthogonal over the distribution of the risk-neutral probabilities given by (3).

The heterogeneous case requires a third argument θ, which appears in the second sum. Here, utilities depend directly on θ through the family (Q_l) of orthogonal polynomials over the distribution of marginal bettors' types (θ). Since higher θ's tend to bet on horses with higher odds R, we center and standardize the distribution on odds for each value of θ. That is, we define the variable $R_\theta = (R - E(R|\theta))/\sqrt{V(R|\theta)}$ and instead of using the 3-tuple (R, p, θ), we use (R_θ, p, θ) as argument of the specification.[23] We chose a family of polynomials (P_k) to be orthogonal over the distribution of R_θ, and once more we normalize them to equal zero when $R = R_M$. As in the homogeneous case, departures from the expected utility case are captured by the family (B_m) of orthogonal polynomials.

These are very flexible specifications, that could approximate any continuous value function to any degree of precision. Even in an expected utility context, we do not impose that the von Neumann–Morgenstern utility be either always concave or always convex; and the normalized fear or ruin may cross the threshold of 1 and fan in or fan out. The price to pay for flexibility is that the monotonicity conditions on V cannot be directly imposed on this form, and will have therefore to be tested ex post on the estimated models. Specifically, V must be increasing in p (which requires that $1 + pK_p \geq 0$) and in R (which requires $1 + (R+1)K_R \geq 0$). Similarly, our single-crossing assumption cannot be translated into simple restrictions on the parameters; rather, we shall check the empirical relevance of the assumption on our preferred estimates. The normalized fear of ruin is

$$\text{NF}(R, p, \theta) = \frac{1 + pK_p(R, p, \theta)}{1 + (R+1)K_R(R, p, \theta)},$$

and Assumption 3 requires that it decrease in θ. We will check this condition on all horses $i < n$ in each race. Finally, the indifference conditions are, in each race and for $i < n$,

$$\log(p_i) + \log(R_i + 1) + K(R_i, p_i, \theta_i)$$
$$= \log(p_{i+1}) + \log(R_{i+1} + 1) + K(R_{i+1}, p_{i+1}, \theta_i), \qquad (14)$$

---

[23]We estimate the conditional expectation and dispersion over the sample (θ_i, R_i), (θ_i, R_{i+1}), excluding the longest shot $i = n$ in each race.

which we need to solve for the probabilities p_i. Define

$$X_i = \frac{1}{R_i + 1} \exp\left(\sum_{j<i}\left(K(R_j, p_j, \theta_j) - K(R_{j+1}, p_{j+1}, \theta_j)\right)\right)$$

with the usual convention $\sum_{\emptyset} = 0$. Since probabilities must sum to 1 in each race, we get

$$p_i = \frac{X_i}{\displaystyle\sum_{j=1}^{n} X_j}. \tag{15}$$

Under EU, this explicitly gives probabilities, since K and therefore the X terms do not depend on **p**. Under NEU, the system of equations has to be solved numerically for each race.[24] Recovering the probabilities allows to compute the log-likelihood function for each race, and therefore for the total sample, as follows:

$$\log L = \sum_c \log p_{w(c)}^c,$$

where $w(c)$ stands for the index of the winning horse in race c. Maximizing this likelihood provides estimates of the coefficients α and β. Notice that this likelihood function, as well as the one in Jullien and Salanié (2000), assumes serial independence between consecutive races. However, the presence of a varying number of horses in common between two races would make accounting for this potential correlation very complex. Another important limitation we share with the existing literature is the assumption that each bettor bets a fixed amount. If bettors were to bet different amounts, but these amounts only depended on, say, individual wealth and were independent of horse and race characteristics, this could be fixed with a re-weighting of the distribution of bettors' types. In that case, each marginal bettor would still be identified, as would his preferences. On the other hand, in the absence of data on individual decisions, we do not have much to say in the case when the amount bet depends on probabilities and odds.

5. RESULTS

We estimate six classes of models, all of which are nested in the general specification given above. Four classes are defined by the distinction between expected utility models and non-expected utility models, and the distinction between homogeneous preferences (whereby all agents have the same attitude towards risk) and heterogeneous preferences (for which we will have to check ex post the validity of the single-crossing assumption). We also introduce two subclasses in the non-expected utility homogeneous class, corresponding to the Yaari (1987) dual model and to the rank-dependent expected utility model (Quiggin (1982), Abdellaoui (2002)). In each class, we still face multiple degrees of freedom, as we can freely vary the number and degrees of the various polynomials. Since many of these models are nested, we shall use the Bayesian information criterion (BIC) to select the best model in each class, and to compare the performances of these models across classes.[25]

---

[24]We used the R package `nleqslv` for that purpose.

[25]We could have used the Akaike information criterion (AIC) instead. These two criteria mainly differ in parsimony. The AIC subtracts twice the number of parameters from the log-likelihood, whereas the BIC

TABLE III

AN OVERVIEW OF MODELS' FIT

Family	BIC	Number of Parameters	$2\log L$
Risk-neutral	$-92,124.3$	0	$-92,124.3$
EU homogeneous	$+49.9$	2	$+70.3$
EU heterogeneous	$+56.7$	3	$+87.3$
Yaari homogeneous	$+29.7$	2	$+50.1$
RDEU homogeneous	$+61.2$	2	$+81.6$
NEU homogeneous	$+73.1$	1	$+83.3$
NEU heterogeneous	idem	idem	idem

Table III provides an overview of our results. We multiplied the value of the log-likelihood by 2 in order to facilitate χ^2 tests. The first line summarizes the benchmark of risk-neutrality, for which the expected return is the same for each horse in any given race. This specification is parameter-free, of course. Each of the following lines describes the best model (as selected by BIC) in a given class. For each class, we list the gain in the value of BIC relative to the risk-neutral benchmark, the number of parameters, and the gain in (twice the) likelihood, compared to the risk-neutral benchmark. It is clear that all selected models perform significantly better in terms of likelihood than the risk-neutral model. What matters, though, is how they perform once the number of parameters they use is taken into account.

This table suggests some surprising conclusions. First, heterogeneity does not seem to play a major role. Indeed, the BIC indicates that, in the comparison between homogeneous EU and heterogeneous EU or between homogeneous NEU and heterogeneous NEU, adding heterogeneity helps little or not at all. Second, non-expected utility seems to matter much more, when compared to expected utility. This is all the more remarkable that while we were able to estimate thousands of models under EU, we only estimated a few hundreds under NEU. From a computational point of view, each NEU model is much more costly to estimate than a EU model since we need to solve a nonlinear system of equations for each race. As a result, Table III underestimates how much NEU overperforms over EU. Third and finally, two homogeneous models dominate: a homogeneous RDEU model with two parameters, and a general homogeneous NEU specification with only one parameter. The next subsections give more information about the estimated shape of preferences in each class. They also discuss whether our preferred models satisfy our theoretical requirements, and in particular the single-crossing condition.

5.1. *Homogeneous Expected Utility*

We start with the simplest specification, in which all bettors are expected utility maximizers with the same attitude towards risk. This can be compared with the "representative agent" model in Jullien and Salanié (2000). The only restriction we have to check ex post is whether utilities are indeed increasing with respect to odds and probabilities. This turns out to hold for all models we estimated.[26] The estimated parameters are significantly

penalizes it by the logarithm of the number of observations. With our 26,525 races, this amounts to 10.2 rather than 2 times the number of parameters. Our experience with AIC is that the number of parameters in the selected models becomes unwieldy, leading to estimates that are sometimes wiggly and do not seem very robust.

[26]For brevity, we relegate to Appendix F of the Supplemental Material the values and standard errors of the parameter estimates (Table S1).

FIGURE 5.—Normalized fear of ruin in the homogeneous expected utility case.

different from zero at a 5% confidence level; the risk-neutral benchmark is thus clearly rejected. The same remark holds for all estimated models.

Figure 5 plots the normalized fear of ruin as a function of odds, along with a 95% confidence band. Remember that NF above 1 reflects risk-aversion. The representative agent appears to be first risk-loving, then slightly risk-averse, then risk-loving again. These estimates are quite different from those reported in Jullien and Salanié (2000). The explanation for this discrepancy lies in their *parametric* approach; they only considered HARA preferences, and they found that within that class, a risk-loving CARA function fit their data best. Our flexible approach shows that assuming a specific functional form is dangerous. For instance, HARA preferences imply a "fanning out" pattern for the fear of ruin: it increases with the odds if and only if it is larger than 1. But our estimated fear of ruin is non-monotonic and crosses the value 1: the data clearly reject the HARA framework.

These preferences indicate that when comparing horses with odds around 6, the representative agent basically only cares about the expected return, while for other comparisons, he behaves in a more risk-loving way, thereby giving a risk premium to relative outsiders. This is consistent with the pattern of expected returns. Figure 8 plots various estimates of the expected return $p(R + 1) - 1$ on each horse as a function of its odds R only. The reference "non-parametric" curve is computed from the raw data on odds and on the identity of the winners. The other curves use the estimated probabilities from our preferred model in each class, instead of the observed frequency of winning. The general picture conforms to the well-known favorite-longshot bias: bettors choose longshots too frequently, so that favorites offer a better expected return. But the nonparametric curve

Quantile	Value of θ
P10	0.375
P25	0.600
P50	0.821
P75	0.938
P90	0.976

appears to flatten for odds between 5 and 10; this helps explain why the NF index of our "representative agent" goes slightly above 1 on this interval.[27]

5.2. *Heterogeneous Expected Utility*

We now turn to heterogeneous expected utility models. Since preferences are now indexed by θ, it is important to examine first the distribution of θ_i, that is, of the types of the bettors who are indifferent between a horse and the next one in a race. The quantiles of θ_i are collected in Table IV. As Figure 6 makes clear, the distribution of θ_i in our sample is much more skewed to the right than the distribution of θ among bettors, which is


FIGURE 6.—Density of the marginal bettors' types θ_i.

---

[27]In an experimental study that examines individual choices between binomial lotteries, Chark, Chew, and Zhong (2016) observed that the favorite-longshot bias is reversed for small probabilities when the expected return is high enough. We do not observe such a reversal; this may be because the expected return in horse races is low, and even negative, while they only proposed lotteries with positive expected returns.

FIGURE 7.—Normalized fear of ruin in the heterogeneous expected utility case.

normalized to be uniform over [0, 1]. There are very few small θ_i's; in fact, since none of our races has more than 12 horses, we cannot observe any θ_i below 1/12. More generally, our observations correspond to the edges of "market share" intervals; and there are many more for outsiders, whose market share by definition is smallest.

The preferred model according to BIC only has four parameters.[28] Figure 7 plots the estimated normalized fear of ruin. The five Pxx solid curves plot $\mathrm{NF}(R, \theta)$ as a function of R for the heterogeneous EU preferences that correspond to the quantiles given in Table IV. The "homogeneous" curve plots the values of NF we found in the homogeneous case.

Figure 5 showed that the estimated preferences under homogeneity were quite complex. Under heterogeneity, these preferences appear as the aggregation of heterogeneous preferences of bettors whose attitude towards risk leads them to self-select into different betting patterns. The individual preferences now are simpler: in particular, for each type, the fear of ruin is almost systematically monotonic with the odds. On the other hand, the fear-of-ruin index often crosses the value 1 that separates risk-aversion from risk-loving. The rejection of HARA preferences seems to be a robust finding, even after we account for aggregation and self-selection. Heterogeneity of preferences also modifies the analysis of the favorite-longshot bias illustrated in Figure 8. Since those bettors who bet on a particular range of horses set the relative price (i.e., the relative odds) for these horses only, allowing for heterogeneity leads to a better fit to the observed expected return. Still, the range 4–8 for odds appears to be special.

An important caveat in our analysis is that the single-crossing property is not always satisfied. This can be seen directly in Figure 7, since some of the curves intersect each other,

---

[28]See Table S2 in Appendix F of the Supplemental Material.

FIGURE 8.—The favorite-longshot bias.

or are even ranked in the "wrong" order. This is especially true for high types and high odds. We computed analytically the derivative of NF with respect to θ for our preferred estimates. The normalized fear of ruin increases significantly in θ for more than half of marginal bettors θ_i, which indicates a robust violation of our single-crossing assumption.[29] One natural interpretation is that the expected utility assumption is simply not supported by our data, leading to spurious violations of the single-crossing assumption. An alternative explanation is that a more complex, multidimensional form of heterogeneity is required to adequately model betting patterns; this calls for further work.[30]

5.3. *Homogeneous Non-expected Utility*

In the NEU case, the results in Table III also support the view that heterogeneity plays at best a minor role. This is why we focus in this part on the homogeneous case.

---

[29]Recall that we selected specifications by rewarding parsimony. It is possible that this exacerbates the violations of the restrictions, in that more flexible specifications would have been more likely to accommodate the restrictions. We chose to err on the side of a more aggressive approach to testing.

[30]Appendix E of the Supplemental Material gives more information on the pattern of these violations.

5.3.1. *Yaari's Dual Model*

Yaari's (1987) model is a natural entry point into the very rich class of non-expected utility models. Comparing the nonlinearity in odds $V(p, R) = pu(R)$ of expected utility to nonlinearity in probabilities $V(p, R) = G(p)(R + 1)$ as in the dual theory nicely frames the question posed by Snowberg and Wolfers (2010): what matters most, preferences or perceptions? By comparing choice patterns across types of bets, they found that a representative "dual" bettor explained the data better than a representative expected utility bettor. Relatedly, Gandhi and Serrano-Padial (2015) assumed that bettors are risk-neutral, but in contrast to our rational expectations equilibrium, they allowed for heterogeneous beliefs. They estimated a model in which roughly 70% of the agents have correct beliefs, while beliefs are noisier for the remaining bettors. Note, however, that they did not impose as much structure as we do, through our single-crossing assumption.

In our richer framework, we can benchmark these two theories using only win bets. As shown in Table III, under homogeneity, expected utility fits the data better than the dual model, although the estimated model allows for quite complex distortions.[31] In a sense, preferences thus seem to matter more than probability distortions in this comparison of opposite models; but we will see that further estimates will lead us to qualify this statement.

We can study these distortions through the lens of the normalized fear-of-ruin index. With a functional $V(p, R) = G(p)(R + 1)$, the indifference equations $V(p_i, R_i) = V(p_{i+1}, R_{i+1})$ identify G up to a multiplicative constant. In principle, this constant could be recovered by imposing $G(1) = 1$; but this point is too far out of the range of observed (risk-neutral) probabilities to be of any use. This should be kept in mind when looking at Figure 9, which plots $G(p)$ as obtained from our preferred model. The figure does suggest that higher probabilities are more likely to be underweighted; this is in accordance with the favorite-longshot bias, and with numerous empirical and experimental observations (see, e.g., Wakker (2010)).

5.3.2. *The RDEU Model*

The RDEU model allows for nonlinearities in both the probability distortion function and the utility from a gain:

$$V(R, p) = G(p)u(R),$$

where, as usual, u is normalized to have $u(-1) = 0$. It clearly offers much more latitude than expected utility; but it still incorporates much more structure than the general NEU specification we shall discuss next. As explained in Section 1.2, both functions G and u jointly determine the normalized fear-of-ruin index:

$$\text{NF}(R, p) = \frac{pG'(p)}{G(p)} \frac{u(R)}{(R + 1)u'(R)};$$

and in the above product, both terms (the elasticity of G with respect to p, and the fear-of-ruin index associated to an expected utility maximizer) are only identified up to a common positive constant. Figure 10 draws this elasticity as a function of the winning probability on our preferred Yaari and RDEU homogeneous estimates. In the Yaari specification, this

---

[31]See Table S3 in Appendix F of the Supplemental Material. Here V is the product of $p(R + 1)$ and of the exponential of a polynomial of degree 3 in p.

FIGURE 9.—The probability distortion function (up to a multiplicative constant) in the Yaari case.

elasticity is uniquely identified; in the RDEU specifications, we can only draw conclusions on its variations. In both cases, the elasticity is rather flat at first, and then decreases for probabilities above 0.3.

5.3.3. *The General NEU Model*

Table III also shows that a general, homogeneous NEU model performs better than a homogeneous RDEU model. The former model also requires only one parameter, which makes it especially parsimonious. This should be contrasted to the findings in Barseghyan et al. (2013) and many other works, which argue that probability distortions are sufficient to explain the choices of deductible in car or auto insurance. These new estimates may point to directions for which the RDEU model is too demanding. Unfortunately, they are difficult to summarize, and we could not find a general feature that would make the NEU model appear superior based on arguments other than statistical fit.[32]

A simple way to contrast the patterns predicted by the different homogeneous models is to plot the normalized fear of ruin $NF(R, p)$ as a function of p for several fixed values of R. This is done in Figure 11 for our four homogeneous specifications, from EU to NEU via Yaari and RDEU. In each panel, we plotted $p \mapsto NF(R, p)$ for the nine deciles of R (P10 to P90) and the nine deciles of p conditional on R. Since NF does not depend on

---

[32]Here is a curious illustration of this similarity. Recall from (13) that all models are the product of $p(R+1)$ and of the exponential of a polynomial K in (R, p). For the RDEU model, K is a linear polynomial in R, *plus* a quadratic polynomial in p (Table S4 in Appendix F of the Supplemental Material); while for the NEU model, K is a polynomial in R, *times* a quadratic polynomial in p (Table S5). This shows both the power of our estimation strategy and the difficulty of understanding analytically the different effects.

FIGURE 10.—The elasticity of the probability distortion function (up to a multiplicative constant in the RDEU case).

p for EU, the first panel plots horizontal segments for each value of R. For Yaari, NF only depends on p and the points align nicely on a single curve. Clearly, that curve is the exception to an inverse U-shaped pattern that appears for EU, RDEU, and NEU. NEU seems to require less departure from smoothness than EU and to accommodate more moderately negative risk-aversion than RDEU, giving the main role to probability distortions.

At this point, our impression is that it is difficult to identify a pattern when comparing these two estimated models. A natural extension would be to test for popular specifications like cumulative prospect theory:

$$V(p, R, \theta) = G^{+}(p, \theta)u^{+}(R, \theta) + G^{-}(1 - p, \theta)u^{-}(\theta).$$

Nonparametric estimation of three separate functions is out of reach, for numerical reasons. We could, of course, resort to parametric specifications, but this is precisely what we have tried to avoid in this paper.

FIGURE 11.—Normalized fear of ruin for four homogeneous models.

CONCLUDING REMARKS

We have argued that it is possible to recover information on the distribution of individual preferences from the sole structure of the equilibrium relationship between prices (in our case, odds) and economic fundamentals (here, probabilities), even in the absence of micro data about individual behavior. We only used four assumptions: agents only care about direct outcomes; the amounts they bet are statistically independent of the lottery they face; they evaluate their decisions using the true probabilities (possibly up to agent-dependent systematic deformations); and a standard single-crossing restriction applies. Then, an equilibrium always exists and is unique. Moreover, the observation of equilibrium patterns overidentifies the distribution of preferences in the population; and underlying assumptions can be tested. Similar ideas have been suggested in several theoretical contributions; but to our knowledge, they have not been taken to data in a systematic way. We provided an empirical investigation of a textbook example, namely, horse races. Our approach could presumably be generalized to more complex frameworks, including insurance and financial markets—but much remains to be done in this direction.

The following three points give a concise summary of our empirical findings:

• It is possible to provide a pretty good description of observed behavior using parsimonious models. While the functional forms we use are highly flexible, the models selected by the BIC (Bayesian information criterion) rely on a very small number of parameters.

• However, obtaining a good empirical fit requires departing not only from standard functional forms such as CARA and CRRA, but from the expected utility framework altogether.

• A simple RDEU model performs as well as the most general NEU specifications. On the other hand, the Yaari dual model is clearly dominated.

Finally, a distinctive feature of our approach is that we consider heterogeneous models as well as homogeneous ones. Our conclusions are mixed on this point. From a methodological standpoint, we show that one-dimensional heterogeneous models of this type can be identified from the data. However, the single-crossing restriction that underlies our approach is not always satisfied by our estimates. Moreover, heterogeneous models do not perform much better than homogeneous ones. It remains to be seen whether this conclusion is linked to the restrictions we impose, most notably the one-dimensionality of heterogeneity. As we showed in Section 2.3, some models with multidimensional heterogeneity can be estimated using strategies that directly generalize our approach. This should be a fruitful program for future research.

APPENDIX: PROOFS

PROOF OF LEMMA 2: From the single-crossing assumption, the set of agents that strictly prefer horse i to horse $j > i$ is an interval containing 0. Similarly, the set of agents that strictly prefer horse i to horse $j < i$ is an interval containing 1. Therefore, the set of agents that strictly prefer horse i to all other horses is an interval. The single-crossing assumption also implies that these intervals are ranked by increasing i; and that the set of agents indifferent between horse i and horse $(i+1)$ is a singleton.   Q.E.D.

PROOF OF PROPOSITION 1: From Definition 1, given a race $(\mathbf{p}, t)$, we have to find a family $\mathbf{R}$ such that, for all $i < n$,

$$V\left(p_i, R_i, (1-t)\sum_{j \leq i}\frac{1}{R_j + 1}\right) = V\left(p_{i+1}, R_{i+1}, (1-t)\sum_{j \leq i}\frac{1}{R_j + 1}\right).$$

From the first-order stochastic dominance assumption, the right-hand side is increasing with R_{i+1}, and is strictly below the left-hand side at $R_{i+1} = R_i$. Moreover, Assumption 2 implies that the right-hand side is strictly above the left-hand side for R_{i+1} high enough. Thus, this equality defines a unique R_{i+1}, such that $R_{i+1} > R_i$. The single-crossing assumption then ensures that the difference between the right-hand side and the left-hand side is growing in θ at the right of $(1-t)\sum_{j \leq i}\frac{1}{R_j+1}$. Since, in addition, $V_R > 0$, this proves that R_{i+1} is an increasing function of R_i, and a non-decreasing function of each R_j, $j < i$. Iterating this remark, we get that each R_{i+1} is an increasing function of R_1. Replacing in (2), we get an equation in R_1 which has at most one solution. Existence follows from the fact that $(R_1, \ldots, R_n)$ forms an increasing sequence, so that by setting R_1 high enough, we get $1/(1-t) > \sum_j 1/(1+R_j)$; and from the fact that when R_1 goes to $-t$, we get $1/(1-t) < \sum_j 1/(1+R_j)$.   Q.E.D.

PROOF OF PROPOSITION 2: If we know the odds, then we know the take and the market shares, from (1) and (2); and we also know the indexes $\theta_i(\mathbf{R})$ of marginal bettors, from (6). There only remains to find a family $\mathbf{p}$ solution to the system

$$\forall\, i < n \quad V(R_i, p_i, \theta_i) = V(R_{i+1}, p_{i+1}, \theta_i).$$

Let us focus on positive probabilities. From the first-order stochastic dominance assumption, the right-hand side is increasing with p_{i+1}, and is strictly above the left-hand side at $p_{i+1} = p_i$. From Assumption 1, it is also strictly below the left-hand side when p_{i+1} goes to zero; therefore, p_{i+1} is uniquely defined, and $p_{i+1} < p_i$. Moreover, p_{i+1} is an increasing function of p_i, and thus of p_1. Finally, p_1 is uniquely determined by $p_1 + \sum_{i<n} p_{i+1} = 1$ (existence follows from checking the cases $p_1 \to 0$ and $p_1 = 1$).   Q.E.D.

PROOF OF PROPOSITION 3: Property (iv) holds, as a simple rewriting of (7). Properties (i)–(iii) follow directly from Assumptions 1–3 and the definition of G in (8). For example, recall that the single-crossing assumption states that, for all $R < R'$ and $\theta < \theta'$,

$$V(R, p, \theta) \leq V(R', p', \theta) \quad \Rightarrow \quad V(R, p, \theta') < V(R', p', \theta').$$

This is equivalent to

$$p' \geq G(R, p, R', \theta) \quad \Rightarrow \quad p' > G(R, p, R', \theta'),$$

and thus G must be decreasing with θ, as required in property (i).   Q.E.D.

PROOF OF PROPOSITION 4: Let us define a function w by

$$w(R, p, \theta) = \frac{G_p}{G_R}(R, p, R', \theta),$$

where, by property (ii) of Proposition 3, the RHS does not depend on R'.

The function w is positive by property (i) in Proposition 3. Now, choose some V whose marginal rate of substitution V_p/V_R is equal to w. We can impose $V_R > 0$ and $V_p > 0$. Since

$$\frac{V_p}{V_R}(R, p, \theta) = \frac{G_p}{G_R}(R, p, R', \theta),$$

there must exist a function $\tilde{G}$ such that

$$G(R, p, R', \theta) = \tilde{G}(V(R, p, \theta), R', \theta).$$

Then (i) implies that $\tilde{G}$ is increasing with V. Moreover, from (iii), it must be the case that $\tilde{G}$ is the inverse of V with respect to p. Let us now prove that V verifies the single-crossing assumption. Assume that $V(R, p, \theta) \leq V(R', p', \theta)$, for $R < R'$. Since $\tilde{G}$ is the inverse of V, we get

$$\tilde{G}(V(R, p, \theta), R', \theta) = G(R, p, R', \theta) \leq p'.$$

Since, from property (i), G is decreasing with θ when $R < R'$, we obtain that for $\theta' > \theta$,

$$\tilde{G}(V(R, p, \theta'), R') = G(R, p, R', \theta') < p'.$$

Since $\tilde{G}$ is the inverse of V, we get $V(R, p, \theta') < V(R', p', \theta')$, so that V verifies the single-crossing assumption, as announced. Finally, since $\tilde{G}$ is the inverse of V, property (iv) can be rewritten as

$$\forall\, i < n \quad V\big(R_i, p_i(\mathbf{R}), \theta_i(\mathbf{R})\big) = V\big(R_{i+1}, p_{i+1}(\mathbf{R}), \theta_i(\mathbf{R})\big),$$

which is exactly the set of equilibrium conditions in Definition 1. Thus, $\mathbf{p}(\mathbf{R})$ characterizes the market equilibria associated to V.   *Q.E.D.*

REFERENCES

ABDELLAOUI, M. (2002): "A Genuine Rank-Dependent Generalization of the Von Neumann–Morgenstern Expected Utility Theorem," *Econometrica*, 70, 717–736. [22]

ANDRIKOGIANNOPOULOU, A., AND F. PAPAKONSTANTINOU (2016): "Heterogeneity in Risk Preferences: Evidence From a Real-World Betting Market," Swiss Finance Institute Research Paper No. 13-53. [4]

ARUOBA, S. B., AND M. S. KEARNEY (2011): "Do Lottery Gamblers Love Risk or Overweight Small Odds?," Working Paper, University of Maryland. [5]

ATHEY, S. (2001): "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," *Econometrica*, 69, 861–889. [10]

——— (2002): "Monotone Comparative Statics Under Uncertainty," *Quarterly Journal of Economics*, 117, 187–223. [10]

AUMANN, R. J., AND M. KURZ (1977): "Power and Taxes," *Econometrica*, 45, 1137–1161. [3,9]

BARBERIS, N. (2013): "The Psychology of Tail Events: Progress and Challenges," *American Economic Review: Papers & Proceedings*, 103, 611–616. [4]

BARSEGHYAN, L., F. MOLINARI, T. O'DONOGHUE, AND J. TEITELBAUM (2013): "The Nature of Risk Preferences: Evidence From Insurance Choices," *American Economic Review*, 103, 2499–2529. [2,4,29]

BARSEGHYAN, L., F. MOLINARI, AND J. TEITELBAUM (2016): "Inference Under Stability of Risk Preferences," *Quantitative Economics*, 7, 367–409. [2,4]

BARSKY, R. B., F. T. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, 112, 537–579. [1,4]

BEETSMA, R. M. W. J., AND P. C. SCHOTMAN (2001): "Measuring Risk Attitudes in a Natural Experiment: Data From the Television Game Show Lingo," *Economic Journal*, 111, 821–848. [2]

BERRY, S. (1994): "Estimating Discrete Choice Models of Product Differentiation," *RAND Journal of Economics*, 25, 242–262. [5]

BERRY, S., AND P. HAILE (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1798. [5]

——— (2016): "Identification in Differentiated Product Markets," *Annual Review of Economics*, 8, 27–52. See also Cowles Foundation Discussion Paper No. 2019. [5]

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 60, 889–917. [5]

BINSWANGER, H. P. (1980): "Attitudes Toward Risk: Experimental Measurement in Rural India," *American Journal of Agricultural Economics*, 62, 395–407. [1]

BONIN, H., T. DOHMEN, A. FALK, D. HUFFMAN, AND U. SUNDE (2007): "Cross-Sectional Earnings Risk and Occupational Sorting: The Role of Risk Attitudes," *Labour Economics*, 14, 926–937. [4]

BRESNAHAN, T. (1987): "Competition and Collusion in the American Automobile Market: The 1955 Price War," *Journal of Industrial Economics*, 35, 457–482. [5]

BROWN, D. J., AND R. L. MATZKIN (1996): "Testable Restrictions on the Equilibrium Manifold," *Econometrica*, 64, 1249–1262. [4]

BRUHIN, A., H. FEHR-DUDA, AND T. EPPER (2010): "Risk and Rationality: Uncovering Heterogeneity in Probability Distortion," *Econometrica*, 78, 1375–1412. [4]

CAMERER, C. F., AND H. K. KUNREUTHER (1989): "Experimental Markets for Insurance," *Journal of Risk and Uncertainty*, 2, 265–300. [4]

CHABI-YO, F., D. P. J. LEISEN, AND E. RENAULT (2014): "Aggregation of Preferences for Skewed Asset Returns," *Journal of Economic Theory*, 154, 453–489. [2]

CHARK, R., S. CHEW, AND S. H. ZHONG (2016): "Individual Longshot Preferences," Working Paper. Available at http://zhongsongfa.weebly.com/research.html. [25]

CHIAPPORI, P.-A., AND M. PAIELLA (2011): "Relative Risk-Aversion Is Constant: Evidence From Panel Data," *Journal of the European Economic Association*, 9, 1021–1052. [1]

CHIAPPORI, P.-A., I. EKELAND, F. KUBLER, AND H. POLEMARCHAKIS (2002): "The Identification of Preferences From Equilibrium Prices Under Uncertainty," *Journal of Economic Theory*, 102, 403–420. [4]

——— (2004): "Testable Implications of General Equilibrium Theory: A Differentiable Approach," *Journal of Mathematical Economics, Special issue in honor of Werner Hildenbrand*, 40, 105–119. [4]

CHIAPPORI, P.-A., A. GANDHI, B. SALANIÉ, AND F. SALANIÉ (2009): "Identifying Preferences Under Risk From Discrete Choices," *American Economic Review, Papers & Proceedings*, 99, 356–362. [5]

CHIAPPORI, P.-A., B. SALANIÉ, F. SALANIÉ, AND A. GANDHI (2019): "Supplement to 'From Aggregate Betting Data to Individual Risk Preferences'," *Econometrica Supplemental Material*, 87, https://doi.org/10.3982/ECTA11165. [3,5]

CHIAPPORI, P.-A., K. SAMPHANTHARAK, S. SCHULHOFER-WOHL, AND R. M. TOWNSEND (2014): "Heterogeneity and Risk Sharing in Village Economies," *Quantitative Economics*, 5, 1–27. [2]

COHEN, A., AND L. EINAV (2007): "Estimating Risk Preferences From Deductible Choice," *American Economic Review*, 97, 745–788. [2,4]

CONLISK, J. (1993): "The Utility of Gambling," *Journal of Risk and Uncertainty*, 6, 255–275. [7]

DOHMEN, T., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP, AND G. WAGNER (2011): "Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences," *Journal of the European Economic Association*, 9, 522–550. [4]

FONCEL, J., AND N. TREICH (2005): "Fear of Ruin," *Journal of Risk and Uncertainty*, 31, 289–300. [9]

GANDHI, A. (2006): "Rational Expectations at the Racetrack: Testing Expected Utility Using Prediction Market Prices," Working Paper, Chicago University. [13]

GANDHI, A., AND R. SERRANO-PADIAL (2015): "Does Belief Heterogeneity Explain Asset Prices: The Case of the Longshot Bias," *Review of Economic Studies*, 82, 156–186. [5,12,28]

GAUTIER, E., AND Y. KITAMURA (2013): "Nonparametric Estimation in Random Coefficients Binary Choice Models," *Econometrica*, 81, 581–607. [5]

GOLLIER, C., AND B. SALANIÉ (2006): "Individual Decisions Under Risk, Risk Sharing and Asset Prices With Regret," Report. [7]

GUISO, L., AND M. PAIELLA (2006): "The Role of Risk-Aversion in Predicting Individual Behavior," in *Competitive Failures in Insurance Markets: Theory and Policy Implications*, ed. by P.-A. Chiappori and C. Gollier. MIT Press and CESifo, 213–250. [1]

HANSEN, L. P. (2007): "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," *American Economic Association Papers and Proceedings*, 97, 1–30. [12]

HANSEN, L. P., AND T. SARGENT (2007): *Robustness*. Princeton University Press. [12]

HARLESS, D., AND C. F. CAMERER (1994): "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica*, 62, 1251–1290. [4]

HEY, J., AND C. ORME (1994): "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, 1291–1326. [1]

JULLIEN, B., AND B. SALANIÉ (2000): "Estimating Preferences Under Risk: The Case of Racetrack Bettors," *Journal of Political Economy*, 108, 503–530. [5,6,14,20,22-24]

KITAMURA, Y., AND J. STOYE (2016): "Nonparametric Analysis of Random Utility Models," Report, Yale. [5]

KÖSZEGI, B., AND M. RABIN (2007): "Reference-Dependent Risk-Attitudes," *American Economic Review*, 97, 1047–1073. [12]

OTTAVIANI, M., AND P. SORENSEN (2010): "Noise, Information and the Favorite-Longshot Bias in Parimutuel Predictions," *American Economic Journal: Microeconomics*, 2, 58–85. [5]

QUIGGIN, J. (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3, 323–343. [22]

SNOWBERG, E., AND J. WOLFERS (2010): "Explaining the Favorite-Longshot Bias: Is It Risk-Love or Misperceptions?" *Journal of Political Economy*, 118, 723–746. [5,6,28]

SUNG, M., AND J. JOHNSON (2008): "Semi-Strong From Efficiency in Horse Race Betting," in *Handbook of Sports and Lottery Markets*, ed. by D. Hausch and W. Ziemba. North Holland. [6]

SYDNOR, J. (2010): "(Over)Insuring Modest Risks," *American Economic Journal: Applied Economics*, 2, 177–199. [4]

WAKKER, P. (2010): *Prospect Theory: for Risk and Ambiguity*. Cambridge University Press. [28]

WEITZMAN, M. (1965): "Utility Analysis and Group Behavior: An Empirical Study," *Journal of Political Economy*, 73, 18–26. [4,6]

WIK, M., T. A. KEBEDE, O. BERGLAND, AND S. T. HOLDEN (2004): "On the Measurement of Risk Aversion From Experimental Data," *Applied Economics*, 36, 2443–2451. [1]

YAARI, M. (1987): "The Dual Theory of Choice Under Risk," *Econometrica*, 55, 95–115. [11,22,28]

ZIEMBA, W. (2008): "Efficiency of Betting Markets," in *Handbook of Sports and Lottery Markets*, ed. by D. Hausch and W. Ziemba. North Holland. [6]

Co-editor Lars Peter Hansen handled this manuscript.