

Estimating Risk Preferences in the Field[†]

LEVON BARSEGHYAN, FRANCESCA MOLINARI, TED O'DONOGHUE,
AND JOSHUA C. TEITELBAUM*

We survey the literature on estimating risk preferences using field data. We concentrate our attention on studies in which risk preferences are the focal object and estimating their structure is the core enterprise. We review a number of models of risk preferences—including both expected utility (EU) theory and non-EU models—that have been estimated using field data, and we highlight issues related to identification and estimation of such models using field data. We then survey the literature, giving separate treatment to research that uses individual-level data (e.g., property-insurance data) and research that uses aggregate data (e.g., betting-market data). We conclude by discussing directions for future research. (JEL C51, D11, D81, D82, D83, G22, I13)

1. Introduction

Risk preferences are integral to modern economics. They are the primary focus of the literature on decision making under uncertainty. They play a central role in insurance and financial economics. The topics of risk sharing and insurance are prominent in development, health, labor,

and public economics, particularly in the study of incentives and social insurance programs. And risk preferences are a major driver in models of consumption, investment, and asset pricing in macroeconomics. While much of the literature is theoretical in nature, deriving qualitative predictions in different environments, there is also a large empirical literature that estimates risk preferences, both their magnitude and their nature.

Most of the literature uses expected utility (EU) theory to model risk preferences. Under EU theory, there are two potential sources of variation in attitudes toward risk: people might differ in (i) their degree of diminishing marginal utility for wealth (their utility curvature), or (ii) their subjective beliefs. Over the years, however, economists have come to recognize additional sources of

*Barseghyan: Cornell University; Molinari: Cornell University; O'Donoghue: Cornell University; Teitelbaum: Georgetown University. We thank Steven Durlauf, four anonymous referees, Pierre-André Chiappori, Liran Einav, and Bernard Salanié for comments, and Heidi Verheggen and Lin Xu for excellent research assistance. Financial support from National Science Foundation grant SES-1031136 is gratefully acknowledged. In addition, Molinari acknowledges financial support from National Science Foundation grant SES-0922330.

[†]Go to <https://doi.org/10.1257/jel.20161148> to visit the article page and view author disclosure statement(s).

variation in attitudes toward risk, and have integrated these into “non-EU” models. The most prominent of these additional sources are (iii) probability distortions (such as in rank-dependent EU) and (iv) reference-dependent utility (as in loss aversion).

Early empirical studies on risk preferences focus on the EU model and rely on data from laboratory experiments (e.g., Preston and Baratta 1948; Yaari 1965); for reviews, see Camerer (1995) and Starmer (2000). Laboratory experiments generated many insights about risk preferences, and most notably demonstrated both substantial heterogeneity in risk preferences and substantial deviations from EU theory. However, the limitations commonly associated with the laboratory setting—e.g., concerns about ecological and external validity—motivated economists to look for suitable data from field settings—i.e., environments in which people’s real-world economic behavior is observable.

As a result, there is a relatively small but growing literature that takes on the difficult task of estimating risk preferences using field data. Our goal in this review is to survey and assess this literature, with a particular emphasis on clarifying the differences among potential sources of variation in risk attitudes and highlighting how one might tease them apart. We concentrate our attention on studies in which risk preferences are the focal object and estimating their structure is the core enterprise. In particular, we generally exclude papers that estimate a structural model of risk preferences, but do not treat the risk preference parameters as the parameters of main interest. Although there are many excellent papers in this category that make important contributions to numerous fields of economics, they are beyond the scope of this review.¹

¹As we explain below, however, we discuss a handful of papers that, although they fall into this category, make

We begin in section 2 with a motivating example designed to address the question of why economists should care about the structure of risk preferences. More and more, economists are engaging in analyses that investigate the quantitative impact of a change in the underlying environment (e.g., a legal reform). In such analyses, risk preferences are often a required input, even if only as part of a broader model. Our example highlights two reasons that the specification of risk preferences matters. First, many quantitative analyses attempt to make out-of-sample predictions for behavior based on the broader model. We demonstrate in our example how different assumptions about risk preferences can lead to different out-of-sample predictions for behavior. Second, many quantitative analyses attempt to reach welfare conclusions. We discuss how different assumptions about risk preferences can lead to different welfare conclusions.

In section 3, we provide a detailed review of several models of risk preferences. Section 3 does not contain an exhaustive list of all models of risk preferences, but rather focuses on those that have been estimated or otherwise studied using field data. We begin with EU theory, and proceed to describe several non-EU models that were originally motivated by experimental evidence, but which subsequently have been studied using field data, including rank-dependent expected utility (RDEU) theory and cumulative prospect theory (CPT). In section 4, we then provide a discussion of identification, and in particular describe what types of data are needed to estimate and distinguish the various models.

In section 5, we discuss research that estimates risk preferences, and sometimes heterogeneity in risk preferences, using individual-level data. We begin with an overview

valuable contributions to the methodology of estimating risk preferences using field data.

of the general approach used throughout the literature. Next, we describe in detail research that estimates risk preferences using data on property-insurance choices. We then briefly discuss studies that use data from television game shows. Lastly, we review a handful of recent papers that analyze data on health-insurance choices. Although health insurance is an important field context, we limit the depth of our coverage because the papers that use health-insurance data do not focus on estimating risk preferences. We believe this is because estimating risk preferences using health-insurance data is especially challenging. Nevertheless, we highlight a few recent papers that address some of these challenges and whose contributions could facilitate future work that focuses on estimating risk preferences.

In section 6, we turn to research that estimates risk preferences, and sometimes heterogeneity in risk preferences, using market-level, or aggregate, data. Once again, we begin with an overview of the general approach of the literature, highlighting how the use of aggregate data naturally requires a stronger set of assumptions in order to identify risk preferences. Next, we describe in detail research that estimates risk preferences using data on betting markets, specifically data on betting in pari-mutuel horse races. We then discuss a select assortment of papers that use macroeconomic data to estimate risk preferences, including data on consumption and investment (asset returns) and on labor supply.

Finally, in section 7 we discuss a number of directions for future research. An under-researched issue is the extent to which risk preferences are stable across contexts. We review the few studies that use field data to investigate this issue, and we highlight the questions left open by these studies. Relatedly, we also discuss the possibility of combining data from laboratory and field settings in order to paint a more complete picture of

risk preferences—and also to gain insight on the question of whether experimental results can be directly applied to make field predictions. Next, we describe the recent literature on using surveys to measure risk perceptions, and we discuss the extent to which survey data might be usefully combined with field data to identify and estimate risk preferences under weaker assumptions. Finally, we discuss the importance of “mental accounting,” by which we mean assumptions about how people translate a complex field context into a set of concrete lotteries to be evaluated. We encourage future research to pay more careful attention to such assumptions.

2. *Motivating Example*

In this section, we present a stylized example designed to motivate why economists should care about the structure of risk preferences. The setting of our example is a hypothetical insurance market. We make a number of strong assumptions—about the setting and the data—that make identification and estimation more straightforward. In later sections, we highlight some of the identification and estimation challenges that economists face in more realistic field settings.

Imagine that there is a continuum of households of measure one who each face the possibility of a loss L that occurs with probability μ . Both L and μ are the same across households, and their values are fixed and known.² To fix ideas, let $L = 10,000$ and $\mu = 0.05$. There is insurance available to the households—full insurance at a price p . Moreover, there is sufficient exogenous price variation (e.g., over time or across various identical subsets of the households; see Einav, Finkelstein, and Cullen 2010 for an example) to non-parametrically identify the market-demand

²Note that by assuming μ is fixed, we are abstracting from moral hazard.

function for full insurance, $Q^F(p)$, which returns the fraction of households willing to purchase full insurance at price p . Panel A of figure 1 depicts one such demand function, namely $Q^F(p) = 2 - 0.001p$. It is a typical demand function—as the price of insurance decreases, the fraction of households willing to purchase it increases. It also reflects aversion to risk—households demand insurance at actuarially unfair prices.

2.1 Out-of-Sample Predictions

Understanding the underlying structure of risk preferences matters for making out-of-sample predictions. Consider a regulatory proposal to require all insurance policies to carry a deductible $d < L$. In order to assess this proposal, we need to know how the demand for insurance would respond to the introduction of the deductible d . The demand function for full insurance $Q^F(p)$ —which we observe—provides, by itself, limited information about the market-demand function for deductible insurance, $Q^D(p)$. However, if we know the underlying model that generates $Q^F(p)$, we can use that model to construct $Q^D(p)$.

Assume for the moment that the underlying model is EU. In addition, assume that (i) the utility function exhibits constant absolute risk aversion (CARA), specifically $u(y) = -\exp(-ry)/r$, where r is the coefficient of absolute risk aversion; and (ii) the slope of the demand function $Q^D(p)$ arises from heterogeneity in r . Given these assumptions, we can recover the population distribution of r , denoted F_{EU} , from the market-demand function $Q^F(p)$. Observe that a household's willingness to pay for full insurance is the z such that³

$$(1) \quad \exp(rz) = \mu \exp(rL) + (1 - \mu).$$

³For an arbitrary utility function, z is defined implicitly by $u(w - z) = \mu u(w - L) + (1 - \mu)u(w)$, where w is the household's status quo wealth.

Equation (1) defines $r^F(z)$ —the coefficient of absolute risk aversion of a household with willingness to pay z for full insurance. A household purchases full insurance when its $r > r^F(p)$, and hence the demand for full insurance satisfies $Q^F(p) = 1 - F_{EU}(r^F(p))$. It follows that, given $Q^F(p)$, we can recover F_{EU} . Panel B of figure 1 displays the F_{EU} that corresponds to the $Q^F(p)$ depicted in panel A.

Given F_{EU} , it is straightforward to construct the demand for deductible insurance $Q^D(p)$. A household's willingness to pay for deductible insurance is the z such that


$$(2) \quad \mu \exp(r(z + d)) + (1 - \mu) \exp(rz) \\ = \mu \exp(rL) + (1 - \mu).$$

Equation (2) defines $r^D(z)$ —the coefficient of absolute risk aversion of a household with willingness to pay z for deductible insurance. A household purchases deductible insurance when its $r > r^D(p)$, and hence the demand for deductible insurance is $Q^D(p) = 1 - F_{EU}(r^D(p))$. Panel C of figure 1 depicts the $Q^D(p)$ that corresponds to the $Q^F(p)$ depicted in panel A, assuming $d = 2,500$. Because deductible insurance provides less coverage than full insurance, naturally $Q^D(p) < Q^F(p)$.


Making a different assumption about the underlying model, however, can lead to different predictions about the level of demand for deductible insurance. Suppose, for example, that the underlying model is the probability distortion model featured in Barseghyan et al. (2013b). The probability distortion model posits that households, instead of weighting outcomes by their objective probabilities, weight outcomes using distorted probabilities.⁴ Under the probability distortion model

⁴In the setting of our example, such probability distortions can emerge from several prominent alternative

Panel A: Demand for full insurance



Panel B: Implied distribution of the risk-aversion coefficient



Panel C: Demand for deductible insurance




Figure 1. Demand for Insurance and Underlying Risk Preferences

(and maintaining the additional assumptions specified above), a household’s willingness to pay for full insurance is the z such that

$$(3) \quad \exp(rz) = \Omega(\mu) \exp(rL) + (1 - \Omega(\mu)),$$

where $\Omega(\mu)$ is the weight on the loss outcome. Suppose that $\Omega(\mu) = \bar{\Omega} > \mu$ is the same across households, and that $\bar{\Omega}$ is known.

models of choice under risk. See section 4.4 for further discussion.

Given $\bar{\Omega}$, we can proceed as before to use the known demand for full insurance $Q^F(p)$ to construct the counterfactual demand for deductible insurance $Q_{\bar{\Omega}}^D(p)$.

Let $F_{\bar{\Omega}}$ denote the distribution of r given the probability distortion model with loss weight $\bar{\Omega}$. We can recover $F_{\bar{\Omega}}$ from the demand for full insurance, $Q^F(p) = 1 - F_{\bar{\Omega}}(r_{\bar{\Omega}}^F(p))$, where $r_{\bar{\Omega}}^F(z)$ is defined by equation (3) with $\Omega(\mu) = \bar{\Omega}$. Panel B of figure 1 displays the $F_{\bar{\Omega}}$ that corresponds to the $Q^F(p)$ depicted in panel A,

assuming $\bar{\Omega} = 0.10$. Given F_{Ω} , we can construct the demand for deductible insurance, $Q_{\Omega}^D(p) = 1 - F_{\Omega}(r_{\Omega}^D(p))$, where $r_{\Omega}^D(z)$ is defined by

$$\begin{aligned} & \bar{\Omega} \exp(r(z+d)) + (1 - \bar{\Omega}) \exp(rz) \\ & = \bar{\Omega} \exp(rL) + (1 - \bar{\Omega}), \end{aligned}$$

the equation that implicitly defines a household's willingness to pay z for deductible insurance. Panel C of figure 1 depicts the $Q_{\Omega}^D(p)$ that corresponds to the $Q^F(p)$ depicted in panel A, assuming $d = 2,500$ and $\bar{\Omega} = 0.10$. Observe that $Q_{\Omega}^D(p) < Q_{EU}^D(p)$.

In short, we see that the two models generate different predictions for the level of demand for deductible insurance. In particular, the EU model predicts a higher level of demand than the probability distortion model. The intuition for this difference follows from the nature of concave utility. Under both models, a concave utility function implies that the concern for reducing risk is stronger the more risk one bears.⁵ Moreover, this effect becomes stronger as the concavity of the utility function increases (i.e., as r gets larger). For a given (observed) demand for actuarially unfair full insurance, the concavity of the utility function is greater under the EU model than under the probability distortion model with $\Omega(\mu) > \mu$,⁶ and thus the implied demand for deductible insurance is greater under the EU model.

⁵Take our example: although the deductible insurance provides 75 percent of the coverage of full insurance, under both modes a household's willingness to pay for the deductible insurance is greater than 75 percent of the willingness to pay for full insurance (see figure 1).

⁶Intuitively, this is because under the EU model a household's aversion to risk (which generates its insurance demand) is driven solely by the concavity of its utility function, whereas under the probability distortion model a household's aversion to risk is driven also by the over-weighting of its distortion function.

2.2 Welfare Analysis

Understanding the underlying structure of risk preferences is also important for welfare analysis. There are two key issues here.

First, when one uses a structural model of preferences to conduct welfare analysis, a misspecified model can yield misleading conclusions. In simple terms, if a misspecified model leads to incorrect out-of-sample predictions for the behavioral impact of a policy change (as in the prior subsection), then of course welfare conclusions will be misleading. However, even if the misspecified model leads to correct predictions for the behavioral impact of a particular policy change, it still can yield misleading conclusions when one uses any welfare analysis that involves interpersonal utility comparisons (e.g., conducting welfare analysis based on sum of surplus). In particular, different underlying models can have different implications for the change in surplus that each person experiences from a particular policy change (even while agreeing on the direction of the change).

Second, economists sometimes conduct welfare analysis without reference to the underlying structure of preferences, using the fact that market demand (or more generally, people's revealed willingness to pay) can be a sufficient statistic for consumer welfare.⁷ For instance, Einav, Finkelstein, and Cullen (2010) propose an approach to empirical welfare analysis in insurance markets that relies only on estimating the demand function. However, this type of welfare analysis is valid only if people's revealed willingness to pay is indeed a sufficient statistic for consumer welfare. The behavioral economics literature has suggested a variety of reasons people's

⁷Such analyses are confined to within-sample welfare analysis, because without an underlying model of preferences, one cannot make out-of-sample behavioral predictions.

behavior might deviate from what maximizes their welfare. Indeed, Baicker, Mullainathan, and Schwartzstein (2015) describe how the standard revealed-preference approach to welfare might fail in the context of health insurance.

The question of whether and, if so, when we should drop the revealed-preference assumption in welfare analysis has been hotly debated—see, in particular, Kőszegi and Rabin (2008), Bernheim (2009), and Chetty (2015). Estimating the underlying structure of preferences can help frame this debate because the more one understands the forces that drive behavior, the better one can assess whether those forces should be given normative weight. To illustrate in the context of risk preferences, suppose we estimate that a probability distortion model (as described in section 2.1) best explains behavior, and suppose we are able to further establish that probability distortions primarily reflect risk misperceptions (i.e., incorrect subjective beliefs). We have then reframed the debate into one about whether we should evaluate welfare using a person’s (incorrect) subjective beliefs or more objective probabilities.

3. Models of Risk Preferences

In this section we describe in detail several models of risk preferences. We begin by reviewing the standard EU model. We then proceed to introduce several alternative models. Our goal is not to provide an exhaustive list, but rather to focus on models of risk preferences that have been prominent in the literature that uses field data to estimate risk preferences.⁸

We start by introducing notation that we use throughout this section.

⁸In the online appendix, we provide further details about these models and illustrate their differences by describing their predictions in three examples.

DEFINITION 1: Let $X \equiv (x_1, \mu_1; x_2, \mu_2; \dots; x_N, \mu_N)$ denote a lottery that yields outcome x_n with probability μ_n , where $\sum_{n=1}^N \mu_n = 1$.

Models of risk preferences describe how a person chooses among lotteries of this form, where we often use \mathbf{X} to denote a choice set. Throughout, we express lottery outcomes in terms of increments added to (or subtracted from) the person’s prior wealth w . In other words, if outcome x_n is realized, then the person will have final wealth $w + x_n$. The probabilities should be taken to be a person’s subjective beliefs. In particular, the models below describe how a person’s subjective beliefs impact his or her choices. The models are silent on the source of those subjective beliefs—we return to this issue in section 5.1.

3.1 Expected Utility

According to EU theory, given a choice set \mathbf{X} , a person will choose the option $X \in \mathbf{X}$ that maximizes

$$EU(X) \equiv \sum_{n=1}^N \mu_n u(w + x_n),$$

where u is a utility function that maps final wealth onto the real line.

Under EU theory, a person’s attitude toward risk is fully captured by her utility function u (and her prior wealth w). In broad terms, a person will be risk averse if u is concave, risk loving if u is convex, and risk neutral if u is linear. More narrowly, one can derive a local measure of absolute or relative risk aversion (or risk lovingness) that characterizes how a person will react locally to choices between lotteries.

Hence, when one estimates an EU model, the main object to estimate is the utility function u . As we shall see, occasionally researchers have taken a nonparametric approach to estimating u , but most often they assume a specific parametric functional

TABLE 1
FUNCTIONAL FORMS USED IN THIS REVIEW

<i>Panel A. Utility functions</i>	
CARA	$u(y) = \begin{cases} -\frac{1}{r} \exp(-ry) & \text{for any } r \neq 0 \\ y & \text{for } r = 0 \end{cases}$
CRRRA	$u(y) = \begin{cases} \frac{1}{1-\rho} y^{1-\rho} & \text{for any } \rho \neq 1 \\ \ln y & \text{for } \rho = 1 \end{cases}$
HARA	$u(y) = \begin{cases} \frac{\gamma}{1-\gamma} \left(\eta + \frac{y}{\gamma}\right)^{1-\gamma} & \text{for any } \gamma \neq 1 \\ \gamma \ln\left(\eta + \frac{y}{\gamma}\right) & \text{for } \gamma = 1 \end{cases}$
NTD	$\hat{u}(\Delta) \equiv \frac{u(w+\Delta)}{u'(w)} - \frac{u(w)}{u'(w)} \cong \Delta - \frac{r}{2} \Delta^2$
<i>Panel B. Probability weighting functions</i>	
Karmarkar (1978)	$\pi(\mu) = \frac{\mu^\gamma}{\mu^\gamma + (1-\mu)^\gamma}$
Tversky and Kahneman (1992)	$\pi(\mu) = \frac{\mu^\gamma}{[\mu^\gamma + (1-\mu)^\gamma]^{1/\gamma}}$
Lattimore, Baker, and Witte (1992)	$\pi(\mu) = \frac{\delta \mu^\gamma}{\delta \mu^\gamma + (1-\mu)^\gamma}$
Prelec (1998)	$\pi(\mu) = \exp(-(-\ln \mu)^\alpha)$
<i>Panel C. Value function</i>	
Tversky and Kahneman (1992)	$v(y) = \begin{cases} y^\alpha & \text{for } y \geq 0, \alpha \in (0, 1) \\ -\lambda(-y)^\beta & \text{for } y < 0, \beta \in (0, 1), \lambda > 1 \end{cases}$

form for u . Perhaps the most common functional forms are the constant absolute risk aversion (CARA), the constant relative risk aversion (CRRRA), and the hyperbolic absolute risk aversion (HARA) families, reported in panel A of table 1.

When one uses the CARA family, one estimates the parameter r , which is the coefficient of absolute risk aversion (higher r means more risk averse). The CARA family

implies a person's prior wealth w is irrelevant to her choices. This is advantageous from the econometrician's viewpoint, because w frequently is unobserved. At the same time, however, this is disadvantageous from the economic theorist's viewpoint, because economists typically believe that people exhibit decreasing absolute risk aversion—i.e., as a person becomes wealthier, she becomes less averse to risk.

When one uses the CRRA family, one estimates the parameter ρ , which is the coefficient of relative risk aversion (higher ρ means more risk averse). The CRRA family has the advantage of implying decreasing absolute risk aversion (among those who are risk averse). However, the CRRA family has the major drawback that it requires prior wealth w as an input. Hence, when researchers use the CRRA family and do not observe prior wealth, they typically either posit some reasonable value for prior wealth (and check robustness for other values) or proxy for wealth using some aspect of the data (e.g., home value).

Finally, when one uses the HARA family, one estimates the parameters η and γ , which together determine the degree of absolute risk aversion $r(y) = (\eta + y/\gamma)^{-1}$. The HARA family has the property that it nests the CARA and CRRA families as special cases, with $\gamma \rightarrow +\infty$ yielding CARA and $\eta = 0$ yielding CRRA.⁹

A third technique is to use an approximation approach; see Cohen and Einav (2007), Barseghyan, Prince, and Teitelbaum (2011), and Barseghyan et al. (2013b). Specifically, if one takes a second-order Taylor approximation of the utility function around prior wealth w and then normalizes by marginal utility evaluated at prior wealth w , one gets

$$(4) \quad \tilde{u}(\Delta) \equiv \frac{u(w + \Delta)}{u'(w)} - \frac{u(w)}{u'(w)} \simeq \Delta - \frac{r}{2} \Delta^2,$$

where $r \equiv -u''(w)/u'(w)$ is local absolute risk aversion. This approximation is accurate when the third- and higher-order derivatives of the utility function u are negligible, at least relative to the increments to wealth that are

relevant in a particular application. As such, we label this approach the negligible third derivative (NTD) approach.

The NTD family is convenient to work with because it does not require prior wealth as an input. However, one must be careful to assess whether the approximation method is appropriate for the particular application under consideration. This will depend on the magnitude of the increments to wealth relative to the estimated degree of risk aversion.¹⁰

3.1.1 Utility Curvature and the Rabin Critique

As it is usually applied—and as it is described above—EU theory is defined for static choices wherein a person faces a single decision problem that involves choosing between lotteries that add to or subtract from her initial wealth.¹¹ For such choices, Rabin (2000) demonstrates that if one assumes that a person uses the same utility function in every choice—an assumption one would want to make if the theory is to have any predictive power—then EU theory is problematic. In particular, Rabin demonstrates that if a person exhibits any noticeable risk aversion over small stakes for a range of initial wealth, then she must exhibit extremely large risk aversion over moderate stakes for that same range of initial wealth, so large as to be clearly counterfactual. For instance, if a person rejects a 50–50 gamble to lose \$10 and win \$10.10 for any initial wealth, then she must also reject a 50–50 gamble to lose \$1,000 and win any positive sum, no matter how large. Because people arguably do exhibit noticeable risk aversion over small stakes, but

¹⁰One obvious concern is that utility must be increasing, which for risk averse individuals (with $r > 0$) holds only for $\Delta < 1/r$.

¹¹Under EU theory, one can equivalently convert this problem into a static choice between lotteries defined over final wealth states.

⁹Some researchers—e.g., Cicchetti and Dubin (1994) and Jullien and Salanié (2000)—assume a simpler HARA specification $u(y) = (\eta + y)^\gamma$. This simplification necessitates restricting γ to lie in the interval $(0, 1]$.

also reasonable risk aversion over moderate stakes, Rabin concludes that EU theory cannot be a good explanation for behavior. This argument is known as the “Rabin critique.”

Motivated by the Rabin critique, parts of the empirical literature have focused on calibrational “rejections” of EU theory, by which they mean a finding of too much utility curvature over small or moderate stakes. In our review of the literature, we describe some examples of such calibrational rejections, where authors conclude that the estimated degree of utility curvature is “too large.” We also attempt to clearly distinguish when EU theory is being rejected for calibrational reasons and when it is being rejected because an alternative model statistically better explains the data.

One possible response to the Rabin critique is that the static EU framework is merely a simplification, as people are in fact solving dynamic life-cycle problems with many decisions taking place over time. If we think of the static EU framework as an “as-if” way of analyzing one of these many decisions, it becomes less clear that we should be applying the same utility function to every decision that the person makes. For instance, for some decisions uncertainty resolves quickly (such as horse race bets or laboratory gambles), while for other decisions uncertainty resolves more slowly (such as property-insurance choices), and thus it is not obvious that the same utility function should apply to all decisions. It is an open question how to think about the comparability of utility functions estimated in different contexts. We return to this issue in section 7.1.

3.2 Rank-Dependent Expected Utility

Rank-dependent expected utility (RDEU) theory emerged from a tradition in psychology of relaxing the feature of EU theory that outcomes are weighted by their probabilities.

In other words, we replace the EU equation with

$$V(X) \equiv \sum_{n=1}^N \omega_n u(w + x_n),$$

where ω_n is a decision weight associated with outcome x_n and may not be equal to a person’s belief μ_n . The original idea was proposed by Edwards (1955, 1962) and popularized in Kahneman and Tversky’s (1979) prospect theory, which assumes $\omega_n = \pi(\mu_n)$. That is, there is an increasing function π —often labeled a probability weighting function—that transforms each probability into a decision weight (still normalizing $\pi(0) = 0$ and $\pi(1) = 1$). With this formulation, however, for any $\pi(\mu) \neq \mu$, it is possible to construct examples in which the theory predicts violations of stochastic dominance—i.e., that people would choose a lottery over another that stochastically dominates it. The source of such predictions is that, unlike under EU theory, when evaluating lotteries, the weights need not sum to one.¹²

Quiggin (1982) proposed a rank-dependent model to solve this problem. Under the rank-dependent approach, when evaluating a lottery $X \equiv (x_1, \mu_1; x_2, \mu_2; \dots, x_N, \mu_N)$, a person first ranks the outcomes from best to worst. Specifically, if the outcomes are ordered such that $x_1 < x_2 < \dots < x_N$, then the weight on outcome n is

$$\omega_n = \begin{cases} \pi(\mu_1) & \text{for } n = 1 \\ \pi(\sum_{j=1}^n \mu_j) - \pi(\sum_{j=1}^{n-1} \mu_j) & \text{for } n \in \{2, \dots, N-1\}, \\ 1 - \pi(\sum_{j=1}^{n-1} \mu_j) & \text{for } n = N \end{cases}$$

¹²For instance, if $\pi(1/3) > 1/3$, then there exists $\bar{y} > 0$ such that the model predicts a person would choose the lottery $(x, 1/3; x - y, 1/3; x - 2y, 1/3)$ over the lottery $(x, 1)$ for all $y \in (0, \bar{y})$.

where π is a probability weighting function. With this approach, when evaluating a lottery, the weights sum to one by construction, and there are no violations of stochastic dominance.¹³

The implications of RDEU theory, of course, depend on the specific probability weighting function that is used. The literature—in large part based on experimental results—has emphasized an inverse-S-shaped probability weighting function: for small μ , $\pi(\mu)$ is concave and has $\pi(\mu) > \mu$, while for large μ , $\pi(\mu)$ is convex and has $\pi(\mu) < \mu$.

Beyond the general inverse-S shape, a number of parameterized functional forms have been proposed in the literature on probability weighting. Some prominent functional forms are reported in panel B of table 1,¹⁴ and depicted in figure 2.¹⁵ Note two features of these functions. First, except for the Karmarkar function, they are not symmetric around $\mu = 1/2$, but rather they typically cross the forty-five-degree line at $\mu < 1/2$. Second, the functions exhibit excess steepness near $\mu = 0$ and $\mu = 1$ —in the sense of $\pi'(\mu) \gg 1$. In fact, in their original discussion of probability weighting, Kahneman and Tversky (1979) instead suggested that probability weighting is discontinuous at

the endpoints, reflecting a notion that as the probability of an event gets small enough, people ignore that possible event. The subsequent literature seems to have introduced the excess steepness near $\mu = 0$ and $\mu = 1$ to eliminate this discontinuity. However, it is unclear how much evidence there is for this excess steepness. As we shall see, in field applications, it is important to assess whether and how low-probability events are incorporated into a person's decision calculations.

3.3 Cumulative Prospect Theory

Kahneman and Tversky's (1979) prospect theory has two key features: probability weighting and loss aversion. As discussed above, probability weighting derived from an older tradition in psychology, and is fully incorporated into RDEU theory. Loss aversion represents a second departure from the EU model: instead of a utility function u defined over final wealth, there is a value function v defined over gains and losses relative to some reference point.

Tversky and Kahneman (1992) propose an improved version of their theory, labeled “cumulative prospect theory” (CPT). CPT requires as an input a reference outcome s , and each outcome is coded as a gain or loss relative to this reference outcome.¹⁶ Consider a lottery $X \equiv (x_1, \mu_1; \dots; x_N, \mu_N)$ and a reference point s , and suppose $x_1 < \dots < x_{\bar{n}-1} \leq s < x_{\bar{n}} < \dots < x_N$. Under CPT, this lottery is evaluated as

$$V(X; r) \equiv \sum_{n=1}^N \omega_n v(x_n - s),$$

¹³While some view rank dependence as merely a technical solution, others attempt to offer intuitive arguments for rank dependence (e.g., Diecidue and Wakker 2001).

¹⁴To be precise, Lattimore, Baker, and Witte (1992) propose the functional form

$$\pi(\mu_i | \mu_{-i}) = \frac{\delta \mu_i^\alpha}{\delta \mu_i^\alpha + \sum_{j \neq i} \mu_j^\alpha}, \quad i = 1, \dots, N,$$

with μ_{-i} denoting the entries of the probability vector μ other than μ_i . For $N = 2$, the above expression coincides with what appears in table 1, and that functional form (used also in the case of $N > 2$) is commonly referred to in the literature as the Lattimore, Baker, and Witte function.

¹⁵Figure 2 closely parallels figure 1 from Prelec (1998). As in that figure, we use $\alpha = 0.65$ for the Prelec function and $\gamma = 0.61$ for the Tversky and Kahnemann function. For the Lattimore, Baker, and Witte function, we use $\delta = 0.77$ and $\gamma = 0.44$, which are estimates from Gonzalez and Wu (1999). For the Karmarkar function, we use $\gamma = 0.50$.

¹⁶The discussion will focus on a reference outcome expressed in increments to wealth, and thus the comparison is x to s . One could equivalently use a reference outcome expressed in final wealth, in which case the comparison would be $(w + x)$ to s —i.e., the value function $v(x - s)$ would be replaced with $v((w + x) - s)$.




Figure 2. Probability Weighting Functions

where the weight on outcome x_n is

$$\omega_n = \begin{cases} \pi^-(\mu_1) & \text{for } n = 1 \\ \pi^-(\sum_{j=1}^n \mu_j) - \pi^-(\sum_{j=1}^{n-1} \mu_j) & \text{for } n \in \{2, \dots, \bar{n} - 1\} \\ \pi^+(\sum_{j=n}^N \mu_j) - \pi^+(\sum_{j=n+1}^N \mu_j) & \text{for } n \in \{\bar{n}, \dots, N - 1\} \\ \pi^+(\mu_N) & \text{for } n = N \end{cases}$$

In this formulation, π^- and π^+ are probability weighting functions applied to the loss and gain events, respectively. Thus, the theory permits differential weighting for gains

and losses.¹⁷ Of course, this differential weighting creates the potential for violations of stochastic dominance.

The value function v is assumed to have three key properties: (i) $v(0) = 0$ and it assigns positive value to gains and negative value to losses; (ii) it is concave over gains and convex over losses (often labeled “diminishing sensitivity”); and (iii) it is steeper in

¹⁷If $\pi^+(\mu) = 1 - \pi^-(1 - \mu)$, then the distinction between π^- and π^+ becomes irrelevant.

the loss domain than in the gain domain (often labeled “loss aversion”).

To estimate a CPT model, one often needs functional form assumptions (although occasionally researchers have attempted more nonparametric approaches). In terms of the probability weighting functions π^- and π^+ , the CPT literature has used the same functional forms as the RDEU literature—indeed, the Tversky and Kahneman function reported in panel B of table 1 was suggested as part of CPT. The value function proposed by Tversky and Kahneman (1992) is reported in panel C of table 1. In that specification, $\alpha \in (0, 1)$ and $\beta \in (0, 1)$ generate diminishing sensitivity in the gain and loss domains, respectively. The parameter $\lambda > 1$ reflects loss aversion, as it implies the negative value generated by a loss is greater than the positive value generated by an equally sized gain. Based on their experimental data, Tversky and Kahneman (1992) suggest that $\lambda = 2.25$, $\alpha = \beta = 0.88$, and for their probability weighting function, $\gamma^- = 0.69$ and $\gamma^+ = 0.61$.

When applying CPT, researchers must specify a reference point, and typically this is done using some external intuitive argument. For instance, in experiments it is typically argued that the reference point should be zero or experimentally endowed wealth. In field settings, researchers often argue for a natural reference point given the setting (e.g., in his recent analysis of tax evasion, Rees-Jones 2018 argues that a zero balance due is a natural reference point). This extra “degree of freedom” in CPT is often seen as a limitation and it has led to various ideas about how to tie down the reference point.

3.4 Expectations-Based Models

A class of “expectations-based” models advances the idea that expectations about outcomes set reference points and influence choices.

Kőszegi and Rabin (2006, 2007) propose a model of loss aversion in which the reference point is taken to be one’s expectations about outcomes. Moreover, because such expectations could involve uncertainty about future outcomes, they extend the model of loss aversion to use a reference lottery instead of a reference outcome.

Specifically, under Kőszegi–Rabin (KR) loss aversion, the utility from choosing lottery $X \equiv (x_n, \mu_n)_{n=1}^N$ given a reference lottery $\tilde{X} \equiv (\tilde{x}_m, \tilde{\mu}_m)_{m=1}^M$ is

$$V(X|\tilde{X}) \equiv \sum_{n=1}^N \sum_{m=1}^M \mu_n \tilde{\mu}_m \times [u(w + x_n) + v(w + x_n | w + \tilde{x}_m)].$$

The function u represents standard “intrinsic” utility defined over final wealth, just as in EU. The function v represents “gain–loss” utility that results from experiencing gains or losses relative to the reference lottery. Gain–loss utility depends on how a realized outcome x_n is compared to all possible outcomes that could have occurred in the reference lottery. For the value function, KR use

$$v(y|\tilde{y}) = \begin{cases} \eta[u(y) - u(\tilde{y})] & \text{if } u(y) > u(\tilde{y}) \\ \eta\lambda[u(y) - u(\tilde{y})] & \text{if } u(y) \leq u(\tilde{y}) \end{cases}$$

In this formulation, the magnitude of gain–loss utility is determined by the intrinsic utility gain or loss relative to consuming the reference point. Moreover, gain–loss utility takes a two-part linear form, where $\eta \geq 0$ captures the importance of gain–loss utility relative to intrinsic utility and $\lambda \geq 1$ captures loss aversion. The model reduces to EU when $\eta = 0$ or $\lambda = 1$.

KR propose that the reference lottery equals recent expectations about outcomes—i.e., if a person expects to face lottery \tilde{X} , then her reference lottery becomes \tilde{X} . However, because situations vary in terms of when a person deliberates about and then

commits to her choices, KR offer multiple solution concepts for the determination of the reference lottery. Here, we focus on two solution concepts that are perhaps most relevant for field data.

DEFINITION 2 (KR–PPE): *Given a choice set \mathbf{X} , a lottery $X \in \mathbf{X}$ is a personal equilibrium if for all $X' \in \mathbf{X}$, $V(X|X) \geq V(X'|X)$, and it is a preferred personal equilibrium if there does not exist another $X' \in \mathbf{X}$ such that X' is a personal equilibrium and $V(X'|X') > V(X|X)$.*

DEFINITION 3 (KR–CPE): *Given a choice set \mathbf{X} , a lottery $X \in \mathbf{X}$ is a choice-acclimating personal equilibrium if for all $X' \in \mathbf{X}$, $V(X|X) \geq V(X'|X)$.*

KR suggest that PPE is appropriate when, faced with a choice set \mathbf{X} , a person thinks about the choice situation, decides on a planned choice $X \in \mathbf{X}$, and then makes that choice shortly before the uncertainty is resolved. An option X is a personal equilibrium if, when a person plans on that option and thus that option determines her reference lottery, it is indeed optimal to make that choice. Among the set of personal equilibria, the PPE is the personal equilibrium that yields the highest “utility.” In terms of field contexts, then, PPE is an appropriate solution concept when a person is able to think about a choice situation for some duration and then make a choice shortly before the uncertainty is resolved. Among those that we discuss in sections 5 and 6, the field context that perhaps best fits this scenario is betting on horse races.

The idea behind CPE is that, when faced with a choice set \mathbf{X} , a person commits to a choice well in advance of the resolution of uncertainty. By the time the uncertainty is resolved, the person will have become accustomed to her choice and hence expect the lottery induced by her choice. Hence, the person chooses the lottery that yields the

largest utility conditional on that lottery being the reference lottery. Two field contexts in which a person commits to a choice well in advance of the resolution of uncertainty are property insurance and health insurance.

When estimating the KR model, one needs to estimate the parameters η and λ along with the utility function $u(y)$. Because the latter is meant to be standard utility over final wealth, as in EU, any of the functional forms for $u(y)$ in table 1 might be used.

Models of “disappointment aversion” also assume that choices are influenced by expectations. The concept of disappointment aversion was proposed by Bell (1985) and further developed by Loomes and Sugden (1986) and Gul (1991). The basic idea is that one is disappointed (or elated) if the realized outcome of a lottery is worse (or better) than expected.

Bell (1985) proposes a variant of disappointment aversion in which disappointment is determined from a comparison of one’s realized utility to one’s EU, and the person accounts for expected disappointment when making a choice. Formally, a lottery $X \equiv (x_n, \mu_n)_{n=1}^N$ is evaluated as

$$V(X) = \sum_{n=1}^N \mu_n u(w + x_n) - \beta \sum_{n=1}^N \mu_n [I(u(w + x_n) < \bar{U}) (\bar{U} - u(w + x_n))],$$

where I is an indicator function and $\bar{U} \equiv \sum_{n=1}^N \mu_n u(w + x_n)$. The first term is the standard EU of lottery X . The second term reflects the expected disutility from disappointment that arises when the realized utility from an outcome is less than the standard EU of the lottery. The parameter β captures the magnitude of disappointment aversion, where the model reduces to EU for $\beta = 0$.¹⁸

¹⁸Bell (1985) further assumes that (i) $u(x) = x$, and (ii) a person might also experience utility from elation when the realized outcome is larger than the expected utility. Even with the latter, however, his model reduces to

Gul (1991) proposes another variant of disappointment aversion in which disappointment is determined from a comparison of one's realized outcome to one's certainty equivalent for the lottery. Formally, a lottery $X \equiv (x_n, \mu_n)_{n=1}^N$ is evaluated as $V(X) = \bar{V}$ such that

$$\bar{V} = \sum_{n=1}^N \mu_n u(w + x_n) - \beta \sum_{n=1}^N \mu_n [I(u(w + x_n) < \bar{V}) (\bar{V} - u(w + x_n))].$$

The z that solves $u(w + z) = \bar{V}$ is one's certainty equivalent for lottery X in this model.

When Bell disappointment aversion is applied to binary lotteries, the model is equivalent to the KR–CPE model. Gul disappointment aversion yields a slightly different model, though the structure is still quite similar. (For equations in the binary insurance case, see section 4.4.) For lotteries with more than two outcomes, the three models are more distinct. For details, see the online appendix.

When estimating models of disappointment aversion, one needs to estimate the parameter β along with the utility function $u(y)$. Because the latter is standard utility over final wealth, as in EU, any of the functional forms for $u(y)$ in table 1 might be used.

4. Model Predictions and Identification

Our goal in this section is to develop intuition for the types of data that may yield point identification of a model's parameters. Point identification obtains when, given any two distinct values for the parameter vector

characterizing the model, one has that the model (when applied with each of these parameter vectors) yields a different predicted distribution for the observable data.¹⁹ The subsections below discuss, for each of the models presented in the prior section, the conditions under which the model's parameters are point identified.

To facilitate our discussion, we focus throughout this section on the example of households purchasing insurance. Specifically, we consider the situation in which a household incurs a loss L with probability μ , but also has the option to purchase insurance against this loss with a deductible $d \geq 0$. The willingness to pay z for such an insurance policy must satisfy the indifference condition

$$(w - z, 1 - \mu; w - z - d, \mu) \\ \sim (w, 1 - \mu; w - L, \mu).$$

Table 2 reports, for each of the models presented in the previous section, the equation implied by this indifference condition—i.e., the equation one would solve to obtain a value for z .

In the empirical applications discussed in this review, typically the observable data are comprised of (i) a discrete choice set (e.g., a set of insurance products); (ii) the characteristics of that choice set (e.g., the premiums associated with each insurance product); and (iii) the option selected from that choice set.²⁰ The willingness to pay z for an insurance product is a useful tool in generating, for such discrete choice sets, the model-implied joint distribution of premiums and optimal choices. Consider, for instance, when the choice set is composed of two options,

the model in the text where β represents the difference between the marginal disutility from disappointment and the marginal utility from elation. Loomes and Sugden (1986) also use this formulation, except they study nonlinear disappointment.

¹⁹See Lewbel (2017) for a thorough discussion of identification in econometrics.

²⁰In some cases, the data also contain some characteristics of the household making the choice

TABLE 2
WILLINGNESS TO PAY (z) FOR INSURANCE WITH DEDUCTIBLE d ,
AGAINST THE POSSIBILITY OF LOSING L WITH PROBABILITY μ

Model	WTP	
EU	$\mu u(w - d - z) + (1 - \mu)u(w - z)$	$= \mu u(w - L) + (1 - \mu)u(w)$
RDEU	$\pi(\mu)u(w - d - z) + (1 - \pi(\mu))u(w - z)$	$= \pi(\mu)u(w - L) + (1 - \pi(\mu))u(w)$
CPT	$\pi^-(\mu)v(-d - z) + (1 - \pi^-(\mu))v(-z)$	$= \pi^-(\mu)v(-L)$
KR-CPE	$\{\mu[1 + \Lambda(1 - \mu)]u(w - d - z) + [1 - \mu[1 + \Lambda(1 - \mu)]]u(w - z)\}$	$= \{\mu[1 + \Lambda(1 - \mu)]u(w - L) + [1 - \mu[1 + \Lambda(1 - \mu)]]u(w)\}$
Bell-DA	$\{\mu[1 + \beta(1 - \mu)]u(w - d - z) + [1 - \mu[1 + \beta(1 - \mu)]]u(w - z)\}$	$= \{\mu[1 + \beta(1 - \mu)]u(w - L) + [1 - \mu[1 + \beta(1 - \mu)]]u(w)\}$
Gul-DA	$\left\{ \frac{(1 + \beta)\mu}{1 + \beta\mu}u(w - d - z) + \left(1 - \frac{(1 + \beta)\mu}{1 + \beta\mu}\right)u(w - z) \right\}$	$= \left\{ \frac{(1 + \beta)\mu}{1 + \beta\mu}u(w - L) + \left(1 - \frac{(1 + \beta)\mu}{1 + \beta\mu}\right)u(w) \right\}$

Note: In KR-CPE, $\Lambda \equiv \eta(\lambda - 1)$.

the option to purchase a particular insurance product and the option to remain uninsured. If z is the willingness to pay for the insurance product, then the model-implied joint distribution of premiums and optimal choices involves choosing the insurance for all premiums less than z and choosing no insurance for all premiums greater than z .

To simplify the exposition (and notation), our discussion of identification assumes having data on a population who share the same wealth w , potential loss L , and loss probability μ . Insurance products are defined by their deductible d —i.e., a choice set will be a set of available deductibles (and possibly also the option not to insure). The data will include, for each household, a premium associated with each available deductible, along with the household's choice. For such data, it is natural to use $z(d)$ to denote the willingness to pay for insurance with deductible d , suppressing the fact that this willingness to

pay also depends on w , L , and μ . Of course, in practice, there is heterogeneity in w , L , and μ (and other household characteristics). If these variables were observable, then all identification arguments below would hold conditional on these observables—indeed, as we'll see the literature often views μ as an observed variable. If these variables are unobserved, then identification can become somewhat more complicated—in this and subsequent sections, we discuss ways to deal with various forms of unobserved heterogeneity.

4.1 Expected Utility

We first consider point identification under EU. From table 2, under EU $z(d)$ satisfies

$$\frac{u(w - z(d) - d) - u(w - L)}{u(w) - u(w - z(d))} = \frac{1 - \mu}{\mu}.$$

When estimating risk preferences, much of the literature assumes a parametric

functional form for u —e.g., CARA, CRRA, or NTD—with a single parameter capturing the magnitude of risk aversion. These functional forms all fall in a class of utility functions that satisfy assumption 1 below.

Denote by $u(y; \phi)$ the parametric utility function, where y is a final wealth state and ϕ is a taste parameter. Assume that u is continuous in both $y > 0$ and $\phi \in \mathbb{R}$, and that $\phi = 0$ if and only if $u(y; \phi) = y$. In addition, maintain the following.

ASSUMPTION 1: (i) $u(y; \phi)$ is increasing in y , and for any $y_0 > y_1 > y_2$, the ratio $R \equiv [u(y_1; \phi) - u(y_2; \phi)] / [u(y_0; \phi) - u(y_1; \phi)]$ is strictly increasing in ϕ . (ii) $\lim_{\phi \rightarrow \infty} R = \infty$ and $\lim_{\phi \rightarrow -\infty} R = 0$.²¹

Assumption 1 naturally associates ϕ with the magnitude of an individual's risk aversion.²² In particular, assumption 1 holds if and only if for any $y_0 > y_1 > y_2$ and $\mu \in (0, 1)$, there exists a $\bar{\phi}$ such that $(y_0, 1 - \mu; y_2, \mu) \succ (y_1, 1)$ for $\phi \in [0, \bar{\phi})$, $(y_0, 1 - \mu; y_2, \mu) \sim (y_1, 1)$ for $\phi = \bar{\phi}$, and $(y_0, 1 - \mu; y_2, \mu) \prec (y_1, 1)$ for $\phi > \bar{\phi}$. In words, whenever a person compares a binary risky lottery to a certain amount in the interval (y_0, y_2) , the person chooses the riskier lottery if her risk aversion ϕ is small enough, and she chooses the certain amount if her risk aversion ϕ is high enough.

Consider the willingness to pay $z(d_0)$ for one specific deductible d_0 . Under EU, for any $u(y; \cdot)$ satisfying assumption 1, each preference parameter ϕ implies a unique $z(d_0)$, where the larger is ϕ (the more risk

averse the household is) the larger is $z(d_0)$ (the more the household is willing to pay for insurance). It follows that a data set in which all households make a choice between the same two options—insurance with deductible d_0 versus no insurance—can be sufficient for point identification of ϕ . In particular, because each ϕ implies a unique $z(d_0)$, each ϕ also implies a unique joint distribution of premiums and optimal choices, and thus there is point identification as long as the data contain sufficient variation in premiums.

The literature most often assumes a parametric functional form for u , not only when estimating EU but also when estimating the alternative models we discuss below. As we have seen, this assumption dramatically simplifies identification, but it is a strong restriction. It would be desirable to be able to trace out the utility function nonparametrically over the relevant support. Doing so can be straightforward if one is willing to assume homogeneous preferences (and has access to data exhibiting the type of variation described above). In practice, however, there is heterogeneity across households, and allowing for this heterogeneity dramatically complicates relaxing the parametric assumption on u . We return to this issue in section 6.2.1.

4.2 Rank-Dependent Expected Utility

We next consider point identification under RDEU. From table 2, under RDEU $z(d)$ satisfies

$$\frac{u(w - z(d) - d) - u(w - L)}{u(w) - u(w - z(d))} = \frac{1 - \pi}{\pi},$$

where we use π in place of $\pi(\mu)$ given our maintained assumption that all households have the same μ .²³

²¹The limit assumption is made merely to guarantee interior solutions in any formal results below. In practice, this assumption is unlikely to be important. NTD does not satisfy this assumption, but result 1, below, holds for NTD as well.

²²Assumption 1 is equivalent to condition (e) in Pratt (1964 theorem 1). As shown there, it is equivalent to assuming that an increase in ϕ corresponds to an increase in the coefficient of absolute risk aversion.

²³In other words, the relationship between π and μ plays no role in the discussion in this section. Hence, the identification results in this section also hold in an EU model when one attempts to estimate both risk aversion

In this case, model predictions depend on both the utility function u and the decision weight π , which complicates identification even when the utility function is parametrically specified as $u(y; \phi)$. As above, consider the willingness to pay $z(d_0)$ for one specific deductible d_0 . Unlike above, it is not the case that each vector of preference parameters (ϕ, π) implies a unique $z(d_0)$. Rather, there is a set of (ϕ, π) pairs consistent with $z(d_0)$. To illustrate, we define $\bar{\pi}(\phi|d, \bar{z})$ to be the required π as a function of ϕ that generates willingness to pay $z(d) = \bar{z}$. We can rearrange the equality above to derive

$$\bar{\pi}(\phi|d, \bar{z}) = \frac{u(w; \phi) - u(w - \bar{z}; \phi)}{[u(w; \phi) - u(w - \bar{z}; \phi)] + [u(w - \bar{z} - d; \phi) - u(w - L; \phi)]}$$

Given this function, any preference-parameter pair $(\phi, \bar{\pi}(\phi|d_0, z(d_0)))$ is consistent with $z(d_0)$. For any u that satisfies assumption 1, $\bar{\pi}(\phi|d_0, z(d_0))$ is decreasing in ϕ , as depicted in figure 3. Intuitively, both an increased risk aversion and an increased decision weight on the loss state imply an increased willingness to pay for insurance. Hence, for a fixed willingness to pay, as risk aversion increases, the decision weight on the loss state must decline in order to keep the willingness to pay unchanged.

Hence, unlike for EU, under RDEU one cannot point identify the vector of preference parameters (ϕ, π) using a data set in which all households make a choice between the same two options. However, it can suffice to observe households choosing between three options. For example, consider households choosing between no insurance, insurance with deductible d_0 , and insurance with

deductible $d_1 > d_0$. The willingness to pay $z(d_1)$ is consistent with another set of (ϕ, π) pairs represented by the curve $\bar{\pi}(\phi|d_1, z(d_1))$, again as depicted in figure 3. As we establish in result 1 below, these two curves cross at only one point, yielding a unique (ϕ, π) pair consistent with both $z(d_0)$ and $z(d_1)$.

RESULT 1: *If $u(y; \phi)$ satisfies assumption 1, then for any $0 < d_0 < d_1 \leq L$ there exists a unique $\bar{\phi}$ such that (i) $\bar{\pi}(\phi|d_0, z(d_0)) = \bar{\pi}(\phi|d_1, z(d_1))$ and (ii) $\bar{\pi}(\phi|d_0, z(d_0)) < \bar{\pi}(\phi|d_1, z(d_1))$ for all $\phi < \bar{\phi}$, and $\bar{\pi}(\phi|d_0, z(d_0)) > \bar{\pi}(\phi|d_1, z(d_1))$ for all $\phi > \bar{\phi}$.*

PROOF:

First, note that because u is increasing in y , $z(d_0) > z(d_1)$ while $z(d_0) + d_0 < z(d_1) + d_1$ (otherwise the household would violate dominance). Define $A(\phi) \equiv u(w; \phi) - u(w - z(d_0); \phi)$, $B(\phi) \equiv u(w - z(d_0) - d_0; \phi) - u(w - L; \phi)$, $A'(\phi) \equiv u(w; \phi) - u(w - z(d_1); \phi)$, and $B'(\phi) \equiv u(w - z(d_1) - d_1; \phi) - u(w - L; \phi)$, in which case $\bar{\pi}(\phi|d_0, z(d_0)) = A(\phi)/[A(\phi) + B(\phi)]$ and $\bar{\pi}(\phi|d_1, z(d_1)) = A'(\phi)/[A'(\phi) + B'(\phi)]$. Hence

$$\begin{aligned} \bar{\pi}(\phi|d_0, z(d_0)) &\geq \bar{\pi}(\phi|d_1, z(d_1)) \\ \Leftrightarrow \frac{A(\phi)}{A(\phi) + B(\phi)} &\geq \frac{A'(\phi)}{A'(\phi) + B'(\phi)} \\ \Leftrightarrow \frac{B'(\phi)}{B(\phi)} &\geq \frac{A'(\phi)}{A(\phi)}. \end{aligned}$$

After algebraic manipulations, assumption 1 yields that $A'(\phi)/A(\phi)$ is a strictly decreasing function of ϕ , where $\lim_{\phi \rightarrow \infty} \frac{A'(\phi)}{A(\phi)} = 0$ and $\lim_{\phi \rightarrow -\infty} \frac{A'(\phi)}{A(\phi)} = 1$. Analogously, assumption 1 yields that $B'(\phi)/B(\phi)$ is a strictly increasing function of ϕ , where $\lim_{\phi \rightarrow \infty} \frac{B'(\phi)}{B(\phi)} = 1$ and $\lim_{\phi \rightarrow -\infty} \frac{B'(\phi)}{B(\phi)} = 0$. The result follows. ■

The key intuition behind result 1 is that probability distortions in isolation yield, for

ϕ and unobserved subjective beliefs (or unobserved risk types), where π would be those subjective beliefs (or risk types). If instead one observes μ and wants to identify the function $\pi(\cdot)$ over some range of values for μ , one needs data as described in the text for all those values of μ .




Figure 3. Identification in the RDEU Model

instance, that the willingness to pay for full insurance is exactly double the willingness to pay to eliminate half the risk (in our example, for $d = L/2$). On the other hand, for standard risk aversion (ϕ), the concern for reducing risk is stronger the more risk one bears, and so the willingness to pay for full insurance is less than double the willingness to pay to eliminate half the risk. Moreover, the larger one's standard risk aversion is, the stronger this asymmetry will be. Hence, the ratio of the willingness to pay to reduce some of the risk relative to the willingness to pay to reduce more of the risk serves to identify how much standard risk aversion is present relative to how much distortion of probabilities is at play. This intuition can be applied in other contexts, as well.

Result 1 shows that a pair $\{z(d_0), z(d_1)\}$ determines a unique pair (ϕ, π) .²⁴ Conversely,

²⁴Note that, for some values of $\{z(d_0), z(d_1)\}$, the unique (ϕ, π) could have $\phi < 0$ (in which case preferences are risk loving).

a pair (ϕ, π) determines a unique pair $\{z(d_0), z(d_1)\}$ and also a unique choice correspondence $C(p_0, p_1)$ that captures the choice from the set $\{d_0, d_1\}$ given premiums $\{p_0, p_1\}$. In turn, $\{z(d_0), z(d_1)\}$ and $C(p_0, p_1)$ determine a unique joint distribution of premiums and optimal choices, and thus there is point identification as long as the data contain sufficient variation in premiums. It follows that a data set in which all households make a choice between the same three options—insurance with deductible d_0 , insurance with deductible d_1 , and no insurance—can be sufficient for point identification of the preference parameters (ϕ, π) , provided there is sufficient variation in premiums.

4.3 EU and RDEU with Unobserved Heterogeneity

The discussion above focuses on identification of the EU or RDEU preference parameters for a population in which everyone has the same preference parameters

(i.e., for a model with homogeneous preferences or a model in which preferences depend on only observable characteristics). It is often important to permit unobserved heterogeneity in preference parameters (as we discuss more in section 5.1.2). How can one point identify the distribution of risk preferences in this case?

Intuitively, the researcher now needs to observe how the population (or at least a large enough representative subsample) allocates itself across the different options within a fixed menu, and moreover the researcher needs to observe this allocation for a sufficiently large set of different fixed menus. For the insurance context that we have been using in this section, this means that the researcher needs to observe data that identify the multinomial distribution of chosen deductibles conditional on premiums, with sufficient variation in premiums (i.e., variation in the choice menu) to identify the distribution of risk preferences.

To illustrate, consider an EU model in the insurance context with a parametric utility function $u(y; \phi)$ that satisfies assumption 1. Suppose that the choice set consists of three deductible options, $d_A > d_B > d_C$, with respective premiums $p_A < p_B < p_C$. Each $p \equiv (p_A, p_B, p_C)$ is a choice menu. One can show that a choice menu p partitions the support of ϕ in three regions— $(-\infty, \phi_{AB}(p))$; $(\phi_{AB}(p), \phi_{BC}(p))$; and $(\phi_{BC}(p), \infty)$ —so that, when facing p , individuals with $\phi < \phi_{AB}(p)$ choose deductible d_A , individuals with $\phi \in (\phi_{AB}(p), \phi_{BC}(p))$ choose deductible d_B , and individuals with $\phi > \phi_{BC}(p)$ choose deductible d_C .²⁵ If we then let F denote the cumulative distribution function of ϕ , for a given p the model predicts that

a fraction $F(\phi_{AB}(p))$ chooses deductible d_A , a fraction $F(\phi_{BC}(p)) - F(\phi_{AB}(p))$ chooses deductible d_B , and a fraction $1 - F(\phi_{BC}(p))$ chooses deductible d_C . The researcher then matches these predictions to the empirical multinomial distribution of chosen deductibles conditional on premiums p , and thereby identifies the value of F at two points. By repeating this exercise for a large number of choice menus, F can be fully learned from the data. In particular, we need variation in p that yields variation in $(\phi_{AB}(p), \phi_{BC}(p))$ that spans the entire support of ϕ , so that the entire distribution F can be identified. We return to discussing this strategy in section 6.2.2.

If instead we consider an RDEU model, there are potentially two sources of unobserved heterogeneity: in standard risk aversion ϕ and in the decision weight π . In principle, a similar approach to what was delineated for the EU model might be applied, building on the insight in figure 3. Suppose that the choice set consists of two deductible options, $d_A > d_B$, with respective premiums $p_A < p_B$, and also the option of no insurance. Each (p_A, p_B) is a choice menu. Given p_A , the curve $\bar{\pi}(\phi|d_A, p_A)$ partitions the (ϕ, π) space into a region that prefers d_A to no insurance and a region that prefers no insurance to d_A . Given p_B , the curve $\bar{\pi}(\phi|d_B, p_B)$ yields an analogous partition for d_B versus no insurance. Finally, given the pair (p_A, p_B) , one can add a third curve that partitions the (ϕ, π) space into a region that prefers d_A to d_B and a region that prefers d_B to d_A . These three partitions can then be combined to generate, as a function of (p_A, p_B) , an overall partition of the (ϕ, π) space into a region that chooses d_A , a region that chooses d_B , and a region that chooses no insurance.²⁶ As above, the researcher can then match these predictions to the empirical multinomial

²⁵Of course, for some p , some of these intervals might be empty (which occurs when p is such that an option is dominated). While the intuition extends in a straightforward way, for simplicity we ignore such p here.

²⁶As above, for some (p_A, p_B) , some of these regions might be empty (which occurs when (p_A, p_B) is such that an option is dominated).

distribution of chosen deductibles conditional on premiums (p_A, p_B) , and thereby obtain some information about the distribution of (ϕ, π) . By repeating this exercise for a sufficiently large set of choice menus, the distribution can be learned from the data.

Nonetheless, to date, point identification of multidimensional heterogeneity in risk preferences has relied mostly upon parametric assumptions about their joint distribution. An interesting departure from this literature appears in Ericson et al. (2018), who build upon and expand the arguments here to provide results on nonparametric identification of unobserved heterogeneity in both risk aversion and risk, where “risk” can be interpreted as subjective beliefs. Their analysis can be related to ours observing that π , as described in this section, can be interpreted as risk and can be learned without connecting it with claim probabilities μ ; observation of μ and variation in μ are only needed to identify π as a function of μ (see footnote 23).

4.4 Distinguishing More Complex Models

Finally, we discuss identification of more complex models that include additional sources of aversion to risk—e.g., KR loss aversion or disappointment aversion (DA)—and perhaps multiple sources at the same time. We organize our discussion in two parts. First, we consider contexts in which households only make choices over binary lotteries. Then we consider contexts in which households make choices over lotteries with more than two outcomes.

If the data contains only choices over binary lotteries $X \equiv (x_1, \mu; x_2, 1 - \mu)$, with $x_1 < x_2$, then all models considered in section 3 (except for CPT and KR-PPE) can be reduced to one in which a household chooses the lottery that maximizes

$$(5) \quad U(X) = \Omega(\mu)u(w + x_1) + (1 - \Omega(\mu))u(w + x_2),$$

where $\Omega(\mu)$ is a *probability distortion* function, as in Barseghyan et al. (2013b). In the online appendix, we show that the function $\Omega(\mu)$ may take one of the following forms:

$$\begin{aligned} \text{Under RDEU:} \quad & \Omega(\mu) = \pi(\mu); \\ \text{Under KR-CPE:} \quad & \Omega(\mu) = \mu(1 + \Lambda(1 - \mu)); \\ \text{Under Bell-DA:} \quad & \Omega(\mu) = \mu(1 + \beta(1 - \mu)); \\ \text{Under Gul-DA} \quad & \Omega(\mu) = (1 + \beta)\mu/(1 + \beta\mu).^{27} \end{aligned}$$

Hence, even if one has data that allows for nonparametric point identification of $\Omega(\mu)$, the underlying models can be distinguished only to the extent that they impose different restrictions on $\Omega(\mu)$. Clearly, KR-CPE and Bell-DA cannot be distinguished from each other. Each of KR-CPE/Bell-DA and Gul-DA impose strong parametric assumptions, and so does RDEU if a parametric functional form is assumed for $\pi(\mu)$ (e.g., any of those in table 1). With these parametric assumptions, the models can be tested to find which best fits the data. However, a more flexible approach would allow for nonparametric probability weighting. If either loss aversion or disappointment aversion are also permitted, identification of the resulting $\Omega(\mu)$ function would not allow one to disentangle these sources of aversion to risk. For example, KR-CPE and RDEU together yield $\Omega(\mu) = \pi(\mu)(1 + \Lambda(1 - \pi(\mu)))$, and similarly for RDEU together with Bell-DA or Gul-DA, so that $\pi(\mu)$ and Λ (or β) cannot be separately identified.

Hence, if the data contain only binary lotteries, without relying strongly on functional form assumptions, the best one can do is focus on identification and estimation of $\Omega(\mu)$. Potential exceptions are represented by models that feature CPT or KR-PPE. Under CPT, $V(X; r)$ does not take the form in

²⁷As we show in the online Appendix, under KR-CPE the parameters λ and η always appear in the lottery evaluation as the product $\Lambda \equiv \eta(\lambda - 1)$, and therefore only this product can be identified.

equation (5) because there is an exogenous reference point and utility varies with whether outcomes are above or below that reference point. Hence, one can potentially separately identify the CPT component of the model by studying how behavior changes as outcomes move above and below the reference point (or, alternatively, by studying how behavior changes as the reference point changes). Under KR–PPE, one cannot define a $U(X)$ independently of the other lotteries in the choice set. Hence, one can potentially separately identify the KR–PPE component of the model by studying how behavior changes as the choice set varies. Of course, one needs data with the right type of variation to pursue either of these approaches.

A frequent assumption in the literature is that subjective beliefs μ coincide with objective expectations (e.g., “objective” claim probabilities), which in turn the econometrician can estimate. However, this assumption may fail in a given application. In that case, when μ is assumed to equal objective expectations, the estimated $\Omega(\mu)$ function captures a mapping Ψ from the estimated objective probabilities to subjective beliefs, thereby yielding another possible source of probability distortions. In fact, as explained before, the relationship between probability distortions and μ plays no role in result 1, hence the same identification results hold for a model where Ω is interpreted as subjective beliefs. In turn, this implies that under the assumption that households’ behavior is governed by RDEU, for example, the estimated $\Omega(\mu)$ function would not necessarily correspond to $\pi(\mu)$, but to $(\pi \circ \Psi)(\mu)$, and in order to identify the entire function $(\pi \circ \Psi)(\cdot)$ on a given range of values for μ , one needs to learn from the data the multinomial distribution of deductibles chosen conditional on premiums and the probability of experiencing a loss for the entire range of μ of interest.

When data contain lotteries with more than two outcomes, more refined inference is possible. For example, as pointed out in Barseghyan et al. (2013a), for lotteries with more than two outcomes, a model with RDEU only and a model with subjective beliefs only, as formalized by the mapping Ψ , generate different predictions. Intuitively, under the mapping Ψ alone, the weight on a particular event is independent of the magnitude of the outcome associated with that event. In contrast, under RDEU alone, the magnitude of the outcome associated with an event impacts the rank ordering of outcomes, and thereby can influence the weight. In addition, KR–CPE and RDEU may generate predictions that are different from those generated under either Bell–DA or Gul–DA. Under either KR–CPE or RDEU, the decision weight assigned to an event depends only on the rank order of the outcome associated with that event. In contrast, under either model of DA, the decision weight assigned to an event depends on whether the outcome is above or below the relevant benchmark that determines disappointment—i.e., the EU of the lottery under Bell–DA, or the certainty equivalent under Gul–DA. Hence, variation in outcomes that does not change the rank order but does change the magnitudes of intermediate outcomes can be used to distinguish these classes of models.

However, KR–CPE and RDEU cannot be separately identified even with choice data on lotteries with more than two outcomes. Indeed, one can show that, for lotteries with any number of outcomes, the combination of KR–CPE and RDEU reduces to an equivalent RDEU model using effective probability weighting $\Omega(\mu) = \pi(\mu)(1 + \Lambda(1 - \pi(\mu)))$. Hence, it is never possible to separately identify KR–CPE and RDEU.²⁸

²⁸Masatlioglu and Raymond (2016) make a similar point using a decision-theoretic approach.

We conclude with two important and related implications of this discussion. First, for models that purport to estimate an RDEU model, the estimated probability weighting functions might in fact reflect a combination of probability weighting and some other phenomenon—e.g., KR–CPE, Bell–DA, or Gul–DA. Second, given this fact, it would seem valuable to take a non-parametric approach to estimating RDEU models, as opposed to restricting attention to the functional forms in table 1. As we summarize in section 5.2, the more recent literature has taken a first step in this direction.

5. Estimation with Individual-Level Data

In this section, we describe research that estimates risk preferences using individual-level data. We begin by discussing the general approach that is broadly used in all of these papers. We then review how that approach has been applied in several different contexts: (i) property insurance, (ii) game shows, and (iii) health insurance.

5.1 The General Approach

Estimating risk preferences from individual-level data typically requires three main steps:

Step 1: Identify a field context in which economic agents make choices between options that involve risk and for which the researcher can obtain data on both the agents' choice sets and the agents' choices.

Step 2: Translate the (typically) rich field-choice environment into a choice between a well-defined set of lotteries (as formalized in definition 1)—so that each model of risk preferences defined in section 3 makes a prediction for which option should be chosen as a function of the taste parameters.

Step 3: Enrich the basic models of risk preferences with some form of heterogeneity, because in practice observationally equivalent agents facing identical choice sets are observed to make different choices.

The following subsections describe steps 2 and 3 in some detail, along with issues that arise in each. It is within these steps that the methodological toolkits of economics need to be applied. Step 1 requires the researcher to (i) identify field contexts that are likely to best reveal risk preferences and (ii) be able to obtain the corresponding data. The existing work using individual-level data has focused on some of the field contexts traditionally emphasized by the theoretical literature (e.g., insurance), for which researchers have been able to obtain sufficiently rich data sets. As more data sets become available, and as economists develop a better understanding of how to approach estimation of risk preferences, we expect the general approach described in this section to be applied more broadly.

5.1.1 Translation into Lotteries

Virtually any field-choice environment involves a rich context and does not present itself in the form of lotteries to which one can immediately apply a model of risk preferences. Hence, to begin the analysis, one needs to first make a series of assumptions that permits translating the rich field context into a choice between a well-defined set of lotteries.

Consideration Sets.—To begin this translation, one needs to make an assumption about what exactly is the set of options under consideration—i.e., the consideration set. In some cases, this appears straightforward. For instance, in papers that study people's deductible choices for property insurance, a natural assumption is that the consideration

set includes each of the possible deductibles. Even here, however, the assumption is less innocuous than it might seem—e.g., in some cases, there are available deductibles that are very rarely chosen, and so one might wonder whether people are actually even considering these possibilities, and one might assume these deductibles are not part of the choice set. Also, most analyses get data from one company, and thus one cannot observe options that a household might have considered purchasing from another insurance company. We return in section 7.4 to discuss when assumptions on the consideration set are likely to be important and when not.

Outcomes.—A closely related assumption the researcher needs to make concerns the set of possible outcomes that might occur. In most field contexts, each option is associated with many possible outcomes. Returning to the example of deductible choices, during the policy period an individual might incur zero losses, one loss, two losses, three losses, and so forth. However, the researcher often restricts the set of possible outcomes. For example, some analyses reviewed below in the context of property insurance assume that households focus only on the possibilities of either zero or one loss, and they ignore the (small-probability event) of multiple losses during a single policy period. We return in section 7.4 to discuss when assumptions on the set of possible outcomes are likely to be important and when not.

Subjective Beliefs.—A particularly important step in the translation from the field context to a set of lotteries concerns what one assumes about agents' subjective beliefs on the likelihood of the possible outcomes. In some field contexts, there exist clear objective probabilities for outcomes—e.g., in games of chance such as state-run lotteries or casino roulette. In most field contexts,

however, objective probabilities either do not exist or are very hard to assess. For such situations, an ideal approach would be to simultaneously estimate both the agents' beliefs and preferences. As we shall see in section 7.3, however, this presents a fundamental identification problem. Hence, the most common approach to date has been to assume "rational expectations," in the sense that agents' subjective beliefs correspond to objective probabilities (often, but not always, as reflected in past or future outcomes). The researcher then either posits a carefully thought-out model of rational expectations formation, or posits a "reduced-form" model, and estimates probabilities over outcomes conditional on the chosen covariates based on realized outcomes and observed covariates. These estimated probabilities are then typically taken as "data," in the sense that they are treated as an observed input when estimating preferences.²⁹

Moral Hazard.—A final important assumption concerns whether probabilities over outcomes depend on one's choices—e.g., whether there is moral hazard in an insurance context. In fact, most analyses that estimate risk preferences assume there is no moral hazard, although a few directly study it. We discuss moral hazard more in section 5.1.3.

5.1.2 Sources of Unobserved Heterogeneity

Given a choice set \mathbf{X} and beliefs μ , each model of risk preferences from section 3 generates an objective function $V(X, \mu, \theta)$ such that the model predicts a person will choose the option $\bar{X} \in \mathbf{X}$ that maximizes $V(\cdot)$. The vector θ denotes a set of taste parameters (which differs depending on the model), and the goal is to estimate θ . When

²⁹In most cases, estimation error in this stage is not accounted for when reporting standard errors for the estimates of preferences.

one applies these models to individual-level data, however, a problem inevitably arises: whereas each model generically predicts a unique optimal choice for any specific \mathbf{X} , μ , and θ , in virtually any data set observationally equivalent agents facing identical choice sets are observed to make different choices.³⁰ Hence, one must enrich the basic models of risk preferences by allowing for unobserved heterogeneity. There are three main approaches adopted in the literature (and some analyses incorporate two or even all three of them). We briefly review each of these approaches below, and refer the reader to Manski (1977) for an early discussion.

Traditionally, the most common approach has been to use a classical *random-utility* model (McFadden 1974) wherein one assumes that agents choose the option $X \in \mathbf{X}$ that maximizes a “total utility” $W(X, \mu, \theta) \equiv V(X, \mu, \theta) + \varepsilon(X)$. Under this approach, the unobserved heterogeneity is specified at the level of utility assigned to option X , and moreover, it is assumed to be additively separable from the underlying utility model reflected in $V(\cdot)$. Hence, $\varepsilon(X)$ is an unobserved term that makes the agent more or less likely to choose option X relative to what the agent’s underlying risk preference $V(X, \mu, \theta)$ would predict. One then specifies a joint distribution for the $\varepsilon(X)$ s for all $X \in \mathbf{X}$, and the random-utility model predicts a probability distribution over the set of possible options. Typically, the $\varepsilon(X)$ s are assumed to be iid with a type I extreme value distribution with scale parameter σ ,³¹

³⁰Indeed, in experiments it is not uncommon for individual subjects to make different choices when presented the same choice situation more than once.

³¹The scale parameter σ is a monotone transformation of the variance of $\varepsilon(X)$, and thus a larger σ means larger variance. In general, one must either normalize the scale of the utility function V or the scale of the choice noise σ . The literature typically uses parametrizations of V that normalize its scale, and thus σ is typically estimated.

in which case the predicted probability of choosing option $X^* \in \mathbf{X}$ is

$$\Pr(X^* | \theta, \sigma) = \frac{\exp(V(X^*, \mu, \theta)/\sigma)}{\sum_{X \in \mathbf{X}} \exp(V(X, \mu, \theta)/\sigma)}.$$

A second approach is to use a *random-preference* model wherein one assumes unobserved heterogeneity in the vector θ of preference parameters (recall that we discussed identification of such models in section 4.3). Under this approach, one maintains that agents choose the option $X \in \mathbf{X}$ that maximizes $V(X, \mu, \theta)$. However, one further assumes that there is unobserved (to the researcher) heterogeneity in θ , and thus the researcher will observe different choices even among observationally equivalent agents facing identical choice sets. Of course, rather than estimate a specific θ , the goal becomes to estimate a distribution $F(\theta)$, which might be done parametrically or nonparametrically. In particular, if $\Theta(X^*)$ is the set of θ such that X^* is the optimal choice, the predicted probability of choosing option $X^* \in \mathbf{X}$ is

$$\Pr(X^* | F) = \int_{\theta \in \Theta(X^*)} dF(\theta).$$

Finally, one can assume unobserved heterogeneity in beliefs. Here, one again maintains that agents choose the option $X \in \mathbf{X}$ that maximizes $V(X, \mu, \theta)$. However, one relaxes the assumption that subjective beliefs μ are “data,” and instead assumes that there is unobserved (to the researcher) heterogeneity in beliefs. One then specifies a distribution for μ —typically with mean equal to objective probabilities—and the model predicts a probability distribution over the set of possible outcomes. In particular, if $\mu(X^*)$ is the set of μ such that X^* is the optimal choice, and if $G(\mu)$ is the distribution of μ , then the predicted probability of choosing $X^* \in \mathbf{X}$ is

$$\Pr(X^* | \theta, G) = \int_{\mu \in \mu(X^*)} dG(\mu).$$

Once one enriches the model by incorporating one or more of these unobserved heterogeneity terms, one then estimates the parameters—both the preference parameters and the unobserved heterogeneity parameters—by making the predicted distribution of choices “match” the empirically observed distribution of choices (using maximum likelihood (ML), Markov chain Monte Carlo (MCMC), or some other econometric technique).

Each of these approaches has advantages and limitations. Incorporating unobserved heterogeneity at the level of preference parameters or beliefs can be appealing because it maintains that choices are driven entirely by the underlying utility model reflected in $V(\cdot)$. However, such models can run into trouble when the choice set includes dominated options—i.e., options that should not be chosen for any values of θ or μ . If in the data a dominated option is chosen with positive probability (something that often happens in practice), the researcher must either introduce some form of random trembles in choice or dismiss the subpopulation that chooses dominated options. For example, Barseghyan, Molinari, and Teitelbaum (2016), which we discuss in section 5.2, use a random-preference model where the curvature of the utility function is constrained to lie in a conservative range, and document that the model is incompatible with the choice of a dominated option made by 13 percent of households in their data. In contrast, a random-utility model predicts that all options in the choice set, even dominated options, should be made with positive probability.

On the other hand, by imposing noise at the utility level in an additively separable way, random-utility models run into a different problem. In many underlying utility models, the preference parameters influence not only the nature of preferences, but also the utility differences between different outcomes. As a simple example, under EU

with a CRRA utility function, an increase in ρ on the one hand makes a person more risk averse, but also decreases the utility differences between any two outcomes. Because, given additive separability, the unobserved heterogeneity term operates on the utility differences, this can generate some perverse features. Apestegui and Ballester (2018) point out one particular implication of this issue: For any fixed σ , the choice probabilities can be nonmonotonic in the preference parameters. For instance, under EU with CRRA (or CARA), increasing ρ (or r) might initially shift choice probabilities more and more toward favoring a safer lottery, but as ρ (or r) gets larger yet the utility differences shrink, and so choice probabilities shift more and more toward choosing at random (and away from favoring the safer lottery).

Of course, when one estimates a random-utility model, one jointly estimates the preference parameters in θ and the extent of unobserved heterogeneity as reflected by σ , and part of the estimation will involve setting the magnitude of σ to be appropriate relative to the level of utility differences given the estimated θ . For instance, under EU with CRRA utility, when observed choices call for the estimate of ρ to be larger, the estimate of σ will just need to be smaller. Hence, this issue seems unlikely to be a major problem when estimating a model of homogeneous preferences. However, it becomes more worrisome when estimating a model with (observed or unobserved) heterogeneity in risk preferences, because typically σ is fixed across individuals and therefore individuals with different risk preferences are differently impacted by the noise term in a systematic way. Under EU with CRRA utility, for instance, those with higher ρ will be more influenced by the noise term (i.e., their choice probabilities will be more random).

We conclude with two final comments about the three approaches above. First, in principle, the unobserved heterogeneity

could reflect either permanent differences across economic agents or random differences within agents across choice situations. In cross-sectional data, one could never distinguish one from the other, but if one observes multiple choices from each agent, one can start to tease these apart. In section 7.1, we discuss in detail the extent to which economic agents exhibit stable risk preferences versus risk preferences that change from context to context. Second, returning again to the comparison with the experimental setting, we note that while the latter permits great control over the issues raised in step 2, experimental analyses also must deal with observationally equivalent subjects facing identical choice sets who make different choices. Hence, experimental analyses also must think through different ways of introducing unobserved heterogeneity.

5.1.3 Moral Hazard and Adverse Selection

Much of the literature estimating risk preferences using individual-level data focuses on the insurance context. In this context, economists have long discussed the problems of moral hazard and adverse selection, and thus we comment on these issues here.³²

In the context of insurance, moral hazard refers to the idea that individuals who have more insurance coverage will have less incentive to take care, and thus are more likely to incur a loss—in other words, people's risk is endogenous to their choices. Such endogenous risk—which in principle could arise in other environments as well—can create problems in two ways for estimation of risk preferences. First, because it is hard to account for it, many analyses assume that it does not exist. For such analyses, if in

fact moral hazard does exist, then estimates of risk preferences might be biased. Second, if moral hazard exists, then there must be a reason why behavior is changing—e.g., in the insurance context, the fact that people take less care when they have more coverage presumably reflects that they get some form of utility from taking less care. If so, then this feature of preferences should be incorporated into the analysis.

Adverse selection refers to the idea that individuals who bear more risk (which is not observable or cannot be priced) are more likely to purchase higher insurance coverage. Adverse selection clearly is a problem for the insurance company. However, it is less of a concern for a researcher trying to estimate risk preferences as long as unobserved heterogeneity in risk can be taken into account, even if this unobserved heterogeneity is correlated with unobserved heterogeneity in risk preferences.³³

We conclude with raising one more issue closely related to unobserved heterogeneity. If there is unobserved heterogeneity in risk preferences, then the researcher needs to be confident that the choice sets faced by the agents are independent of those risk preferences. In particular, risk preferences are estimated by investigating how agents react to changes in choice sets, and thus our estimates would be biased if a change in choice sets were correlated with a change in risk preferences. In the insurance context, for instance, if an insurance company can easily get a sense of a person's risk preferences, the company might be able to increase profits by adjusting the insurance pricing in reaction to those preferences. In practice, most insurance pricing is not done in this way—often due to heavy regulation—and thus this issue is perhaps less of a problem.

³²There is a large literature in economics that tests for the presence of moral hazard and adverse selection in insurance markets. We briefly summarize the empirical findings of that literature in the online appendix.

³³For a leading example, see Cohen and Einav (2007), which is discussed in detail in section 5.2.

5.2 Property-Insurance Data

Insurance choices are a natural domain in which to estimate risk preferences. Due to difficulties in obtaining data prior to the 2000s, there are relatively few papers that use individual-level insurance data to carry out this task.

The first paper to use individual-level data on insurance choices to estimate risk preferences is Cicchetti and Dubin (1994). They analyze data from Mountain Bell on roughly 10,000 residential telephone customers in Colorado in 1990. The choice of interest is whether customers purchased inside-wire insurance. This insurance cost roughly \$0.45 per month, and protected against telephone-wire problems inside one's residence. Without the insurance, in the event of a problem, the household would need to pay the service cost to fix the problem, which averaged about \$55. The probability of experiencing a problem was roughly 0.5 percent (see below), and thus the expected benefit of the insurance was roughly \$0.275. Hence, purchasing this insurance for \$0.45 is a clear sign of aversion to risk. Cicchetti and Dubin set out to get a more precise understanding of this aversion to risk.

In order to translate the field context into the domain of preferences (a choice between lotteries, as discussed in section 5.1), Cicchetti and Dubin (1994) assume that households effectively choose between the following two lotteries:

$$(-p, 1) \text{ versus } (-L, \mu; 0, 1 - \mu).$$

The first lottery is that associated with the choice to purchase the insurance, where p is the premium charged for the insurance. The second lottery is that associated with the choice not to insure, where μ is the probability of experiencing a problem within any given month, and L is the expected service

cost in the event of a problem.³⁴ The authors estimate μ using data on historical (1982–86) trouble calls, where they divide the data into nine zones, and they take the observed ratio of trouble calls to customers in each zone to be the probability of a problem for customers in that zone. Then, in their empirical analysis, they treat p , L , and μ as data.

Cicchetti and Dubin (1994) estimate an EU model with a HARA utility specification. For prior wealth—a required input in the HARA specification—they use a measure of monthly income generated from census data. They also allow the curvature of the utility function to depend on a household's average monthly bill.³⁵ In addition to estimating an EU model, they also estimate an RDEU model with a parametrized form of overweighting of probabilities—the one in Lattimore, Baker, and Witte (1992).

They estimate these models by ML, where they use a McFadden (1974) random-utility specification to account for observationally equivalent households making heterogeneous choices. From their estimates, Cicchetti and Dubin (1994) conclude that “the overall pattern of results is remarkably consistent with expected utility theory” (p. 183). In particular, they find virtually no evidence of overweighting of probabilities—i.e., RDEU does no better than EU—and they argue that 78 percent of households have an estimated utility function consistent with EU. Finally, they note that, for the average household, the estimated degree of absolute risk aversion

³⁴This formulation implicitly assumes that households expect, at most, one problem in any given month. As we mentioned in section 5.1, this type of assumption is common in the literature.

³⁵They motivate this assumption based on a reduced-form finding that households with larger average monthly bills are more likely to purchase inside-wire insurance. While this assumption will capture that feature in their structural estimation of risk preferences, it is not clear that there is a good primitive justification for it.

is relatively small and yields a willingness to pay for the insurance virtually identical to the expected benefit from the insurance.

The Cicchetti and Dubin (1994) paper has a number of limitations. Perhaps most important is a data limitation: their data contain very little variation in p , L , and μ .³⁶ Hence, their estimation primarily identifies the impact of monthly income and average monthly bill on the (local) degree of absolute risk aversion. Moreover, as we discussed in section 4, the limited variation in p , L , and μ creates a major impediment to separately identifying both utility curvature and probability weighting, and thus their RDEU estimates depend heavily on their functional-form assumptions for utility and probability weighting. As such, one should be cautious in interpreting their estimates for the degree of probability overweighting.

Finally, upon closer look, it is not so clear how EU in fact fares. For the 22 percent of households with an estimated utility function inconsistent with EU, the inconsistency is in the form of a decreasing utility function (i.e., less is better), which is clearly counterfactual. Among the 78 percent of households with an estimated utility function consistent with EU, nearly half of them (37 percent of all households) are estimated to be risk loving. Moreover, in addition to permitting choices to depend on risk preferences, Cicchetti and Dubin (1994) also include a direct preference for insurance meant to reflect “shifts in tastes and preferences that may be attributable to the insured versus uninsured states” (perhaps due to a belief that households with insurance will get priority service relative to households without insurance). In the estimated model, this direct preference is estimated to be quite large and appears to be the

³⁶The paper seems to suggest that, in the estimation, a single value of p and L is used for all households. While there is some variation in μ across the nine zones, its range of [0.32 percent, 0.74 percent] is very tight.

main factor explaining the choice to purchase insurance for many households.³⁷ Despite these limitations, the Cicchetti and Dubin paper was the first of its kind in the domain of insurance, and inspired later papers that were able to work with better data and more sophisticated models.

A key contribution in this direction is made by Cohen and Einav (2007), which provides a much more sophisticated analysis, although the authors limit themselves to the EU framework. They use individual-level data from an insurance provider in Israel, and analyze deductible choices among households who purchased one particular form of auto insurance (similar to comprehensive automobile insurance in the United States). In their data, all households have purchased the insurance, and the decision of interest is which of four deductible options they chose. The data contain the full menu of premium–deductible combinations offered to each household, along with that household’s chosen deductible. In addition, they observe actual claims made by these households during the policy year.

Since virtually all households (98.9 percent) chose one of the two lowest deductibles, Cohen and Einav (2007) limit attention to the choice between those two.³⁸ Hence, they translate the field context into the domain of preferences—i.e., a choice between lotteries—by assuming that households were

³⁷Relative to the Rabin (2000) critique, this direct preference for insurance term is how they were able to explain the purchase of insurance even at reasonable levels of risk aversion. Yet, there has been some controversy on whether the fact itself that individuals purchase insurance for such small risk at a price that is twice the expected loss can be considered consistent with EU theory, especially in light of the relatively small estimates of absolute risk aversion; see, e.g., Camerer (2000), Rabin and Thaler (2001), and Fehr-Duda and Epper (2012).

³⁸Cohen and Einav (2007) do not drop households that chose one of the two high deductibles, but rather they proceed as if those households had chosen the highest remaining deductible.

effectively choosing between the following two lotteries:

$$(-p_l - d_l, \mu; -p_l, 1 - \mu)$$

versus

$$(-p_h - d_h, \mu; -p_h, 1 - \mu).$$

The first lottery is that associated with the choice of the lower deductible, and the second lottery is that associated with the choice of the higher deductible. For each choice $k \in \{l, h\}$, p_k is the premium charged for the insurance, d_k is the deductible that must be paid by the household in the event of a loss, and μ is the probability of experiencing a loss during the policy period. The insurance is typically priced such that lower deductibles are actuarially unfair—i.e., a risk-neutral household would choose the higher deductible d_h , and d_l would be optimal if the household is risk averse enough.³⁹

This formulation reflects three simplifying assumptions. First, it assumes that each household's deductible choice does not influence its claim probability μ —i.e., it assumes that there is no moral hazard with respect to the deductible choice. Second, it assumes that every possible loss is larger than the larger deductible d_h . Cohen and Einav (2007) explicitly discuss these assumptions, and argue that they are supported in their data. The third simplifying assumption is that households make decisions while considering a policy period that is infinitesimally small, thereby implicitly yielding that households ignore the possibility of incurring more than one claim during the policy period. But since multiple claims are rare (in their data, only 2.70 percent of households experience more than one claim over the one-year policy period), this assumption is probably not

³⁹Under their benchmark estimates of claim rates, this is true for 98.7 percent of households (p. 752).

restrictive (as we discussed in section 5.1, similar assumptions are made in subsequent analyses).⁴⁰

Cohen and Einav (2007) then estimate an EU model with an NTD specification for the utility function.⁴¹ Importantly, they permit both observed heterogeneity (i.e., dependence on household observables) and unobserved heterogeneity in both the degree of (absolute) risk aversion r and the likelihood of a claim μ . Indeed, the main goal of their analysis is to assess (i) the extent of such heterogeneity, (ii) the relative importance of heterogeneity in risk versus in risk preferences, and (iii) the correlation between the unobserved elements of the heterogeneity. Formally, they assume that claims are generated by a Poisson process with Poisson claim rate λ , and that r and λ have a bivariate log-normal distribution with

$$\begin{pmatrix} \ln \lambda_i \\ \ln r_i \end{pmatrix} \overset{iid}{\sim} N \left(\begin{bmatrix} x_i' \beta_\lambda \\ x_i' \beta_r \end{bmatrix}, \begin{bmatrix} \sigma_\lambda^2 & \rho \sigma_\lambda \sigma_r \\ \rho \sigma_\lambda \sigma_r & \sigma_r^2 \end{bmatrix} \right).$$

In the above expression, σ_λ^2 is the variance of λ , σ_r^2 is the variance of r , and ρ is their correlation.

Cohen and Einav (2007) estimate this model using an MCMC approach. The data contain several key features that permit them to parametrically identify the unobserved heterogeneity in risk preferences and in risk. First, they observe the full distribution of (the number of) household claims over the course of the one-year policy period. This distribution permits them to directly

⁴⁰Although Cohen and Einav's (2007) approach effectively assumes that households ignore the possibility of multiple claims during the policy period, their estimation still uses the full distribution of number of claims over the policy period, because this distribution is needed to identify the variance and correlations of unobserved heterogeneity in risk (as we discuss below).

⁴¹In fact, their utility equation is slightly different from the one in section 3 because they consider a policy period that is infinitesimally small.

estimate the mean and variance of risk (for a fixed set of observables) without any reference to premium–deductible menus or to choices.⁴² Second, different households in the data are offered different menus of (p_l, d_l) versus (p_h, d_h) .⁴³ Seeing how households respond to different menus (while knowing the mean and variance of risk) permits identification of the mean and variance of risk aversion. Finally, seeing how these responses differ between households with different claims experiences (i.e., households with a different number of actual claims experienced during the one-year policy period) permits identification of the correlation between the unobserved components of risk and risk preferences.

Cohen and Einav's (2007) estimation yields several interesting conclusions. First, in terms of the degree of risk aversion, mean risk aversion is quite large ($r = 0.0019$), but in fact the distribution is quite skewed and the median is more reasonable ($r = 0.0000073$). In other words, in the estimated model, the roughly 18 percent of households that choose the low deductible are primarily explained by very high risk aversion. Second, they find more unobserved heterogeneity in risk aversion than in risk, and moreover, given their estimates, the unobserved heterogeneity in risk preferences is more important for profits and pricing. Finally, they find a strong positive correlation between unobserved heterogeneity in risk preferences and unobserved heterogeneity in risk. This is driven by the fact that as the observed number of claims

increases, the proportion of households that choose the low deductible increases rapidly. They are cautious in concluding too much from this positive correlation, because the unobservables that drive a correlation between risk aversion and risk might be very context specific.

Sydnor (2010) uses similar data to also study the implications of EU for insurance deductible choices. However, he does not pursue an estimation of preferences, but rather a calibration approach in the spirit of Rabin (2000). Nonetheless, this paper is instructive for those who want to estimate risk preferences using insurance data. Sydnor uses individual-level data from a large home insurance company, from which he obtained a random sample of 50,000 home insurance policyholders in a single state in a single year. At this company, each policyholder must choose one of four deductibles—\$1,000, \$500, \$250, or \$100—where a lower deductible implies a higher premium. Much like Cohen and Einav (2007), Sydnor observes the full menu of premium–deductible combinations offered to each household, along with that household's chosen deductible. In addition, he observes actual claims made by these households during the policy year, from which he can derive claim rates.

Sydnor (2010) translates this field context into the domain of preferences exactly as in Cohen and Einav (2007), by assuming that households are effectively choosing between four lotteries of the form

$$(-p_d - d, \mu; -p_d, 1 - \mu),$$

where p_d is the premium charged for the insurance, d is the deductible that must be paid by the household in the event of a loss, and μ is the probability of experiencing a loss during the policy period. This formulation reflects three simplifying assumptions similar to those in Cohen and Einav (2007).

⁴²Note that if they only observed whether or not households made any claims, they could only estimate the mean claim rate. The variance of claim rates is identified by comparing the likelihood of at least one claim to the likelihoods of at least two or more claims.

⁴³This variation is not purely idiosyncratic. For instance, for two-thirds of households, the menu satisfies $d_h = 0.5p_h$, $d_l = 0.3p_h$, and $p_l = 1.06p_h$, where the variation across households comes entirely from variation in p_h . But there is some additional variation, which Cohen and Einav (2007) discuss in some detail.

In his main analysis, Sydnor (2010) considers an EU model with CRRA utility, and he calibrates the degree of relative risk aversion (ρ) that households would need to have to explain their choices. More precisely, for the 6,268 customers who were new at the sample firm in the sample year—and who thus were more likely to have actively chosen their deductible—he focuses on the 3,791 customers who chose either the \$500 or the \$250 deductible.⁴⁴ Choosing a deductible smaller than the maximum (\$1,000) reveals an aversion to risk, and Sydnor calibrates for each household a lower bound on how risk averse that household must be. In his baseline calibration, he assigns to each household a claim rate equal to the average claim rate among those who chose that deductible, and he assumes a prior wealth of \$1 million. He demonstrates that, for the vast majority of the households that chose deductibles smaller than \$1,000, this specification implies implausibly large risk aversion. He further demonstrates that this finding is robust to assuming a variety of values for prior wealth and to assigning to each household a fitted claim rate (in much the same way as described below for Barseghyan et al. 2013b). Based on these calibrations, Sydnor concludes that EU is not a good explanation of people's deductible choices, and he then discusses several potential alternative explanations, including risk misperceptions, probability weighting, and KR loss aversion.

Barseghyan et al. (2013b), BMOT henceforth, use similar data to estimate a probability distortion model (as described in section 4.4). In other words, they consider a model that, in the insurance-deductible context, nests EU, RDEU, KR loss aversion, Bell or Gul disappointment aversion, and risk misperceptions, and they investigate which combination best explains the variation in

⁴⁴Because only 3 of the 6,268 new customers chose the \$100 deductible, he does not analyze that group.

the data.⁴⁵ Their data come from a large insurance company that sells multiple lines of coverage. The full data set comprises yearly information on more than 400,000 households who held auto or home policies between 1998 and 2006. For their main analysis, BMOT use a core data set of 4,170 households who were new customers at the firm in 2005 or 2006 and who purchased both home and auto insurance. The authors focus on households' initial choices. In this group, every household made a deductible choice for three coverages: home all-perils insurance, auto collision insurance, and auto comprehensive insurance. For each choice, the data contain the full menu of premium-deductible combinations offered to each household, along with that household's chosen deductible.⁴⁶ In addition, the data contain the history of claims for all households.

The translation of the field domain into the domain of lotteries is exactly as in Cohen and Einav (2007) and Sydnor (2010), except that BMOT make one further assumption: they assume that households bracket the three choices narrowly in the sense that, while each household has fixed risk preferences that apply to all three choices, the household treats these choices as three independent decisions of the form studied by Cohen and Einav (2007) and Sydnor (2010). It is worth highlighting that

⁴⁵As discussed in section 4, BMOT can only identify an overall probability distortion, and cannot break it down into these possible underlying forces. Nonetheless, this approach identifies the extent to which these underlying forces together might help better explain deductible choices (relative to EU).

⁴⁶The available deductible options were the same for all households. For home, the options were \$100, \$250, \$500, \$1,000, \$2,500, and \$5,000; for collision, the options were \$100, \$200, \$250, \$500, and \$1,000; and for comprehensive, the options were \$50, \$100, \$200, \$250, \$500, and \$1,000. In order to make it plausible that every potential claim is larger than the largest deductible, BMOT exclude the two highest deductibles in home (chosen by only 1.6 percent of households); households that chose such deductibles are kept in the analysis and treated as if they chose a \$1,000 deductible (this approach is analogous to that in Cohen and Einav 2007).

the assumption of narrow bracketing is in fact implicit in Cohen and Einav (2007) and Sydnor (2010), in the sense that they assume that people have risk preferences that apply to their single observed deductible choice without reference to all the other risk choices that households are making.⁴⁷

BMOT estimate a probability distortion model. Applied to the context of choosing deductibles, the utility from choosing deductible d is

$$\Omega(\mu)u(w - p_d - d) + (1 - \Omega(\mu))u(w - p_d),$$

where μ is the (objective) probability of a claim, u is a utility function, and Ω is a probability distortion function. As a preliminary step, BMOT use the full sample and the full history of claims to estimate, for each coverage, a Poisson panel regression with random effects. They then use the output from these claim rate regressions to assign to each household a fitted probability of a claim for each coverage. In the estimation of preferences, these claim probabilities are treated as data.⁴⁸

In their benchmark analysis, BMOT estimate a model of homogeneous preferences, i.e., where all households have the same utility function u and the same probability distortion function Ω . They assume an NTD specification for u (see table 1; the analysis is then repeated for CARA and CRRA utility

functions), but they take a nonparametric approach to Ω . Specifically, they consider three nonparametric approaches: (i) a polynomial expansion of $\ln\Omega(\mu)$, (ii) a polynomial expansion of $\Omega(\mu)$, and (iii) a cubic spline. All three approaches yield the same conclusion: large probability distortions characterized by substantial overweighting of claim probabilities and only mild insensitivity to probability changes. Moreover, while there also is statistically significant curvature in u , economically the lion's share of households' observed aversion to risk is attributed to probability distortions. This result represents a more convincing demonstration than Sydnor (2010) of the limitations of EU in explaining deductible choices. Whereas Sydnor (2010) merely provides calibration arguments against EU based on the required curvature of the utility function being "too large," BMOT find that a model with large probability distortions better explains the data without imposing any restrictions on the magnitude of curvature in the utility function. BMOT compare special cases of models yielding probability distortions, including KR loss aversion and Gul disappointment aversion (see section 3.4 for details on these models), with the nonparametric estimate of $\Omega(\mu)$, to assess whether these models are consistent with the empirical evidence. The results show that neither KR loss aversion alone nor Gul disappointment aversion alone can explain the estimated probability distortions, thereby suggesting a key role for probability weighting in individuals' deductible choices.

BMOT next expand their analysis to incorporate both observed and unobserved heterogeneity in both the curvature of the utility function and in the magnitude of the probability distortions. When allowing for unobserved heterogeneity, BMOT's econometric model takes the form of a mixed logit (with parametric assumptions on the distribution of the unobserved heterogeneity terms), which they estimate via MCMC.

⁴⁷In their benchmark estimation, the only advantage (relative to Cohen and Einav 2007) of having three choices per household is that it yields increased variation in claim probabilities and prices represented in the data. When BMOT incorporate observed and unobserved heterogeneity into the analysis, this feature will then add a restriction that each household must have the same risk preferences across all three coverages.

⁴⁸They demonstrate that the main results are robust to instead assigning to each household a fitted distribution of claim probabilities for each coverage. As such, the approach in BMOT is roughly equivalent to that in Cohen and Einav (2007), except that they do not permit the unobserved heterogeneity in risk to be correlated with the unobserved heterogeneity in risk preferences.

The mean estimated probability distortions in a model that allows for observed heterogeneity only, unobserved heterogeneity only, or both, are nearly identical to each other and to the estimated probability distortions in a model with homogeneous preferences. Hence, whether BMOT assume preferences are homogeneous or allow for observed or unobserved heterogeneity, their main message is the same. By contrast, the estimated degree of standard risk aversion is somewhat sensitive to the modeling approach. The variance estimates for the unobserved heterogeneity terms suggest a substantial presence of unobserved heterogeneity (though smaller than in Cohen and Einav 2007); unobserved heterogeneity in r and in probability distortions are estimated to be negatively correlated.

Barseghyan, Molinari, and Teitelbaum (2016) use the same data and focus on the same probability distortion model, where they assume CARA utility. However, they allow for completely unrestricted heterogeneity in r and Ω . Specifically, they assume that each household has its own preferences $(r_i, \Omega_i(\cdot))$, and they make the key identifying assumption that these preferences do not change as the household makes risky choices across the three distinct but closely related property insurance contexts (we return to a discussion of this stability assumption in section 7.1). With this assumption, they can partially identify (provide bounds for) each household's risk preferences.

In the probability distortion model, a household's deductible choice in coverage j implies a lower and an upper bound on its probability distortion as a function of r_i and μ_j :

$$(6) \quad LB_i(r_i, \mu_j) \leq \Omega_i(\mu_j) \leq UB_i(r_i, \mu_j).$$

Because the authors observe three choices per household, typically with different μ_j , they obtain three pairs of bounds—or

intervals—per household. Much as discussed in section 5.1.2, however, some households make choices that cannot be rationalized by any $(r_i, \Omega_i(\mu_j))$, and thus the authors exclude 13 percent of the households from the analysis.⁴⁹

Without further assumptions, the restrictions in equation (6) are all that can be learned from the data and the model. However, the various models reviewed in section 4.4 that lead to probability distortions (with the exception of $\Omega_i(\cdot)$ representing subjective beliefs) yield that the function $\Omega_i(\cdot)$ is monotone in μ . The authors therefore focus on exploring the identifying power of this shape restriction to tighten the bounds, and then explore four additional shape restrictions: quadraticity, linearity, unit slope, and zero intercept.⁵⁰ Under each restriction, for each household, the authors search for the minimal r_i such that the household's intervals satisfy the shape restriction. They then analyze that household's lower and upper bounds in equation (6) given that minimal r_i . Of course, for each successive restriction, more households are excluded because there is no $(r_i, \Omega_i(\mu_j))$ that can rationalize their choices while satisfying that shape restriction.

The authors first analyze the distribution of the minimal r_i s under each shape restriction. They find that this distribution is skewed to the right in each case, and that the median is zero under each nondegenerate shape restriction (i.e., each shape restriction other than zero intercept). Indeed, they find that the vast majority of households—between 70 and 80 percent—can be rationalized by

⁴⁹In fact, while some of these households could be rationalized with a very large r_i , the authors also choose to exclude households that cannot be rationalized with an r_i that lies in a reasonable range. Specifically, for reasons they explain in section 5 of their paper, they require $r_i < 0.0108$.

⁵⁰The shape restrictions are cumulative. That is, quadraticity requires that $\Omega(\cdot)$ is both monotone and quadratic, i.e., $\Omega(\mu) = a + b\mu + c\mu^2$, $b \geq 0$, and $c \geq -b/2$; linearity requires $\Omega(\mu) = a + b\mu$, $b \geq 0$; unit slope requires $\Omega(\mu) = a + \mu$; and zero intercept requires $\Omega(\mu) = \mu$.

a model with linear utility given monotone, quadratic, or even linear probability distortions. By contrast, they find that fewer than 40 percent can be rationalized by a model with concave utility but no probability distortions. Next, they use the intervals to perform kernel regressions of the lower and upper bounds on $\Omega(\mu)$ as a function of μ (under each restriction including only households that can be rationalized under that restriction). Under each nondegenerate shape restriction, the results evince a distortion function that substantially overweights small probabilities. Under monotonicity, for instance, the midpoints of the estimated bounds imply $\Omega(0.02) = 0.11$, $\Omega(0.05) = 0.17$, and $\Omega(0.10) = 0.25$, and even the estimated lower bounds imply $\Omega(0.02) = 0.07$, $\Omega(0.05) = 0.12$, and $\Omega(0.10) = 0.15$. The results are very similar under quadraticity and linearity. Moreover, these results are largely consistent with the results in BMOT, which were obtained under more stringent functional form assumptions.

Barseghyan, Molinari, and Teitelbaum (2016) further show how one can utilize the intervals to obtain a point estimate $\Omega(\cdot)$ for the representative household without relying on parametric assumptions about the nature of unobserved heterogeneity. Specifically, they build an estimator based on finding the $\Omega(\cdot)$ that comes closest to rationalizing the behavior of households with monotone $\Omega(\cdot)$. The estimated $\Omega(\cdot)$ can rationalize all three choices of nearly one in five households whose behavior is consistent with monotone probability distortions. Moreover, the residual deviation between the households' intervals and the estimated $\Omega(\cdot)$ allows one to uncover the lower bound on the degree of heterogeneity in probability distortions among households. An interesting next step for future research is to generalize this approach to broader models of decision making under uncertainty.

5.3 Game Shows

Several papers estimate risk preferences using data from television game shows. Game shows provide an attractive setting for estimating EU and non-EU models. Contestants “are presented with well-defined choices where the stakes are real and sizable, and the tasks are repeated in the same manner from contestant to contestant” (Andersen et al. 2008, p. 361). Field settings rarely have such desirable characteristics. That said, a game show is conducted in a peculiar environment that raises concerns about ecological and external validity and observer and selection bias. For this reason, we provide only a brief discussion of the pioneers of the literature and references to a selection of other studies.

The pioneers of the game-show literature include Gertner (1993), who uses data from *Card Sharks*, Metrick (1995), who uses data from *Jeopardy!*, and Beetsma and Schotman (2001), who use data from the Dutch show *Lingo*.

Gertner (1993) assumes that *Card Sharks* contestants are EU maximizers with CARA utility and pursues two methods for estimating the lower bound on the coefficients of absolute risk aversion, r . First, Gertner looks solely at the contestants' bets in the final round of the game and obtains a lower bound on each contestant's r by assuming that a contestant who wagers her entire stake is risk neutral and a contestant who wagers the minimum amount (half of her stake) wishes to bet exactly that. He reports an average lower bound of 0.000310. Second, Gertner looks at the contestants' bets in all rounds and compares the sample distribution of outcomes with the distribution of outcomes if a contestant plays the optimal strategy for a risk-neutral contestant. By revealed preference, an “average” contestant prefers the former distribution to the latter, and so the degree of absolute risk

aversion that would make a contestant indifferent between the two distributions provides an estimate of the average lower bound on r . Under this method, Gertner reports an average lower bound of 0.0000711.⁵¹

Metrick (1995) also models *Jeopardy!* contestants as EU maximizers with CARA utility. In the pertinent part of the paper, Metrick looks only at the bets of the leaders in the final round of “runaway” games, i.e., games in which the leader is so far ahead entering the final round that she can guarantee a win by betting a sufficiently small amount. From the first-order necessary condition of the leader’s utility maximization problem, Metrick derives an expression for the leader’s subjective probability of answering correctly given her bet and coefficient of absolute risk aversion r . This expression is exactly the form of a logit regression. He then obtains an ML estimate of r for the “representative” contestant (i.e., he obtains the r that maximizes the likelihood of the observed sample of bets and correct/incorrect answers). He reports a statistically insignificant point estimate of 0.000066.⁵²

Beetsma and Schotman (2001) consider EU maximization with both CRRA and CARA utility functions. They analyze the decision at the start of a round in the *Lingo*

finals to stop or continue play of the game. They first show that, under fairly weak conditions, the stop/play decision amounts to a choice between receiving the current stake x with certainty and a lottery in which they receive $2x$ with probability p and a with probability $1 - p$, where a represents the option value of coming back in the next show.⁵³ They then specify a probit model for the stop/play decision and estimate the degree of risk aversion. For the CRRA specification, they estimate coefficients of relative risk aversion ranging from 0.42 (assuming zero wealth) to 6.99 (assuming wealth of 50,000 Dutch guilders) to 13.08 (assuming wealth of 100,000 Dutch guilders). For the CARA specification, they estimate a coefficient of absolute risk aversion of 0.12.

Most of the subsequent papers in the game-show literature use data from the United States and international versions of *Deal or No Deal*. For surveys, see Andersen et al. (2008), Post et al. (2008), and Hartley, Lanot, and Walker (2014). Two exceptions are Fullenkamp, Tenorio, and Battalio (2003), who use data from *Hoosier Millionaire*, and Hartley, Lanot, and Walker (2014), who use data from *Who Wants to be a Millionaire*? Like the pioneering studies, many of the subsequent papers work with EU models with CRRA and/or CARA utility (e.g., Fullenkamp, Tenorio, and Battalio 2003; Andersen et al. 2008; Deck, Lee, and Reyes 2008; Conte et al. 2012; Hartley, Lanot, and Walker 2014). Others go beyond EU and consider RDEU, prospect theory, or other non-EU models (e.g., Botti et al. 2008; Post et al. 2008; Mulino et al. 2009; de Roos and Sarafidis 2010; Bombardini and Trebbi 2012).

⁵¹In the second part of the paper, Gertner (1993) presents evidence that contestants’ bets are inconsistent with EU maximization. More specifically, he finds that contestants’ bets exhibit sensitivity to accumulated winnings in the current round (stakes), but not to accumulated winnings in previous rounds (wealth).

⁵²In another part of the paper, Metrick (1995) looks at the bets of the first- and second-place contestants entering the final round of games in which these contestants can mostly ignore the actions of the third contestant. In this part, however, Metrick does not focus on estimating the contestants’ risk preferences. Rather, he focuses on testing whether the contestants play “empirical-best-responses,” i.e., best-responses to the empirical frequency of strategies played by their opponents in his sample of similar games. He finds that while first-place contestants generally play empirical-best-responses, second-place contestants do not.

⁵³Under the rules of the game, if contestants lose, they can come back in the next show unless the current show is their third.

5.4 Health-Insurance Data

Recently, a growing empirical literature has made use of structural models of decision making under uncertainty to answer a variety of important questions about health-insurance markets. These include, *inter alia*, the welfare cost of inertia and imperfect information in health insurance and the implications of adverse selection and moral hazard in these markets.

However, the health insurance field context has not been used for estimation of risk preferences as the object of fundamental interest. We suspect the reason is that estimation of risk preferences when using health-insurance choices is more challenging than it is when using property-insurance choices for at least three reasons. First, the set of outcomes associated with each lottery is significantly more complex; for example, health expenses have (essentially) a continuous distribution. Second, individuals may care about more than mere monetary costs of care; for example, the quality of their life is differentially impacted by different health outcomes. The researcher then either needs to model the utility from each possible health status directly, or to monetize health status. Third, moral hazard is likely to be a larger concern in this context. In particular, if individuals' choices of health-insurance plans are subject to selection on moral hazard, this selection cannot be ignored in the estimation of preferences.⁵⁴

Even when the ultimate target of the analysis is not estimation of risk preferences, the literature studying health-insurance choices needs to contend with the difficulties listed

above. One approach has been to not model directly individuals' risk aversion and simplify the problem by posing a reduced-form equation for individuals' valuation of health insurance plans; see for example Starc (2014). Another approach has been to exogenously impose risk preferences, as opposed to estimating risk preferences within a larger model. See, for example, Einav, Finkelstein, and Schrimpf (2010), who use external data to calibrate the values for risk aversion (as well as other parameters that would only be identified via functional form assumptions in their model).

Because of this, we substantially limit the depth of our coverage of this literature. Nevertheless, we highlight a few recent papers that address some of these challenges and estimate a fully specified EU model with CARA preferences. In certain cases, these papers also allow for preference heterogeneity across individuals. We expect that these contributions could facilitate future work that focuses on estimating risk preferences.

Handel (2013) leverages a major change to insurance provision that occurred at a large firm to quantify the welfare cost of consumer inertia in health insurance markets, and to study policies that could mitigate this inertia. In order to make progress on this important question, Handel assumes away two of the challenges listed above: he assumes that there is no moral hazard and that individuals base their choice of insurance plan only on the monetary costs of each option. However, Handel uses a very careful approach to account for complex lotteries over outcomes.

Handel (2013) observes households choosing among three "preferred provider organization" (PPO) plans, denoted PPO₂₅₀, PPO₅₀₀ and PPO₁₂₀₀. These PPO plans differ among each other in their premiums and cost-sharing characteristics (e.g., deductible, coinsurance, and out-of-pocket maximums). These characteristics, in turn, determine the

⁵⁴Yet another issue is that health-insurance plans can be quite complex, and thus there might be issues of consumer competence with these plans that bias the estimation of risk preferences. Indeed, Bhargava, Loewenstein, and Sydnor (2017) find evidence of large numbers of employees at one company choosing health-insurance plans that are dominated by other available options.

mapping from total medical expenditures to employee out-of-pocket expenditures.

The households in the sample are observed making insurance-plan choices at multiple points in time. Specifically, at time $t = 0$, they are observed making an “active” choice, because the company changed the plans offered (their premiums and cost-sharing characteristics) and the households had to choose a plan among the new offerings. At times $t = 1$ and $t = 2$, households are observed making (potentially) “passive” choices, in the sense that they may simply continue with the plan chosen at time $t = 0$ without reassessing their options. Of notice is the fact that, over time, the PPO₂₅₀ (the plan that yields the more comprehensive coverage) becomes substantially more expensive, while the PPO₅₀₀ (a plan that yields a lower coverage) becomes substantially less expensive; see figure 2 in Handel (2013). As such, choices that were optimal at time $t = 0$ need not be optimal at times $t = 1, 2$.

Handel’s interest is in quantifying the effect of the *inertia* displayed by individuals who do not change their plan over time. To achieve this goal, he sets up a random-utility model of insurance choice, similar to the one in BMOT described above. One important difference is that, in the case of health insurance, the lottery is defined over out-of-pocket expenditures, rather than over binary outcomes. The distribution of out-of-pocket expenditures is assumed known by the households. In practice, it is estimated using medical predictive software, household characteristics, and ex post claim realizations. Specifically, for each individual and open enrollment period, the author uses the past year of diagnoses, drugs, and expenses, along with age and gender, to predict mean total medical expenditures for the upcoming year using the Johns Hopkins ACG Case-Mix software package. He then incorporates medically relevant metrics such as type and duration of specific conditions, as

well as comorbidities. This is done for four distinct types of expenditures: (i) pharmacy; (ii) mental health; (iii) physician office visit; and (iv) hospital, outpatient, and all other. Individuals are then grouped into cells determined by the mean predicted future utilization. For each expenditure type and risk cell, Handel estimates a spending distribution for the upcoming year based on ex post observed cost realizations, combining the marginal distributions across expenditure categories into joint distributions using empirical correlations and copula methods. He then maps individual total expense projections into the family out-of-pocket expense projections, taking into account family-level plan characteristics. This yields the distribution of out-of-pocket expenditures to be used in the choice model.

The effect of inertia is modeled as an additive “consumption” term that is positive for the chosen option in the previous year, yielding

$$c_0 = w_0 - p_0 - e_0,$$

$$c_1 = w_1 - p_1 - e_1 + \eta I(IC_1 = IC_0),$$

where c_t denotes consumption, w_t denotes wealth, p_t denotes price of insurance, e_t denotes out-of-pocket expenditure, each at time $t \in \{0, 1\}$, IC_t denotes insurance plan chosen at time $t \in \{0, 1\}$, and $I(\cdot)$ is the indicator function of the event in parenthesis. In addition, for households with a very high medical expenditure risk, the consumption equation includes an additional dummy variable that puts additional weight on the most comprehensive choice, P_{250} , to capture the idea that essentially all of these households have chosen P_{250} .

The so-specified consumption enters a CARA utility function, which in turn households use to evaluate the three PPO options (by integrating it against the distribution of out-of-pocket expenditure). Then

households choose the option with the highest EU.

The coefficient of risk aversion r is modeled as a random coefficient (following a normal distribution with mean specified as a linear function of observable characteristics), while the coefficient of inertia η is modeled as a linear function of family status and demographics. The resulting econometrics model is a mixed logit and is estimated via MCMC. Unobserved heterogeneity in risk, however, is not allowed for.

The estimates of the coefficient of absolute risk aversion are similar in magnitude to the ones in BMOT, and they lie in $[1.9, 3.25] \cdot 10^{-4}$. The intercept in the inertia function, η , is estimated to be large in magnitude with values of \$1,729 for single employees and of \$2,480 for employees who cover at least one dependent. Because inertia is linked to multiple dimensions of observable heterogeneity, Handel also reports the estimated mean and variance of inertia: The mean total money left on the table per employee due to inertia is \$2,032 with a standard deviation of \$446.

A natural extension of Handel's (2013) model is to study additional frictions that may affect an individual's choice of insurance plans. For example, Handel and Kolstad (2015) study the effect of information frictions and hassle costs, measured via a survey administered to individuals for whom health-insurance information is available. The survey elicits individuals' information about available medical providers/treatments and perceived time and hassle costs to learn the characteristics of a high-deductible health-insurance plan (which is cheaper than the other options). The additional friction measures appear to be important predictors of choices and to significantly decrease risk preference estimates.

Einav et al. (2013) address the difficulties associated with evaluation of outcomes in the health insurance context, as well as the presence of moral hazard. In particular,

moral hazard and its impact on selection is the focus of their paper.

The authors begin their analysis with the key observation that households' utilization rate of medical services may depend significantly on the characteristics of their health insurance coverage. They propose the following simple model, which builds on a model in Cardon and Hendel (2001). Each period is divided into two subperiods. In the second subperiod, households take insurance coverage as given, and their utility function is assumed separable in health and money

$$u(m; \lambda, \omega) = h(m - \lambda; \omega) + y(m),$$

where m is the monetary value of the health care utilization chosen; λ is the monetary value of the health shock; $y(m)$ is the residual income, which is decreasing in m at a rate that depends on the health-insurance coverage; and ω is a parameter that captures households' responsiveness to the price of medical utilization.

The authors assume that $h(m - \lambda; \omega) = m - \lambda - \frac{1}{2\omega}(m - \lambda)^2$ and that, under insurance contract j , the marginal cost of health care is c_j (i.e., $y'(m) = -c_j$). It follows that the optimal amount of utilization is

$$(7) \quad m_j^*(\lambda) = \lambda + \omega(1 - c_j).$$

While these assumptions are quite strong, they play a key role in keeping the model tractable. In particular, through equation (7) the authors assure that the health shock contributes to utilization additively separably from the contribution of moral hazard.

Denoting by $F_\lambda(\cdot)$ the distribution of the health shock, the authors aim at identifying $F_\lambda(\cdot)$ and ω . To argue that these functionals are point identified, the authors rely heavily on (i) the fact that incremental utilization, which is observed, does not depend on risk aversion, but on ω ; and (ii) that the data contains a plausibly exogenous change in

the entire menu of health coverage options. First, the authors consider a counter-factual (but ideal) data setting in which one observes each household's entire distribution of medical expenditure for two different coverages (denoted I and II), call them $G_i^I(m)$ and $G_i^{II}(m)$. Assuming that the distribution of observed health shocks $F_\lambda(\cdot)$ is the same under both coverages, using the fact that expenditure should equal $\lambda + \omega(1 - c_j)$, and denoting by $E_{G_i^k}(m)$ the expected value of utilization under distribution $G_i^k(m)$, $k \in \{I, II\}$, it immediately follows that the difference $(E_{G_i^I}(m) - E_{G_i^{II}}(m))/(c_I - c_{II})$ uncovers the parameter ω . The distribution $F_{i,\lambda}(\cdot)$ for each individual i can also be learned. This is because the authors observe panel data. If individuals within this panel are observed for a sufficiently long period under the different coverages, and if $F_{i,\lambda}(\cdot)$ is time invariant, then $F_{i,\lambda}(\cdot)$ can be learned recalling that $\lambda_{it} = m_{it} - \omega_i(1 - c_i)$, with t denoting time period.

Risk aversion is introduced in the analysis by monetizing the second-period utility $u(m_j^*(\lambda); \lambda, \omega)$ and using this object as the argument of a CARA utility function with coefficient of risk aversion r . Hence, the EU from coverage j in this model is

$$EU_j^\omega = -\frac{1}{r} \int \exp(-u(m_j^*(\lambda); \lambda, \omega)) dF_\lambda(\lambda).$$

An individual specific coefficient of standard risk aversion r_i can then be identified if individuals face a continuous option set.

The authors argue that the data they have comes sufficiently close to this ideal scenario. Specifically, their data hails from Alcoa, Inc., a large multinational producer of aluminum and related products, and it comprises health insurance choices and medical care utilization of their US-based workers (and their dependents). The analysis in this paper focuses mostly on data from the years 2003 and 2004. The data are very rich, document-

ing health insurance options and choices, claim information, demographic characteristics of households, etc., as well as a summary proxy of individuals' health-care utilization based on predictive medical software.

For this paper, the distinctive feature of the data is that in 2004 Alcoa introduced a new set of health insurance PPO options, which were phased in gradually to different employees based on their union affiliation, since new benefits could only be introduced when an existing union contract expired. The staggered timing in the transition from one set of insurance options to another provides a plausibly exogenous source of variation in coverage (mimicking the possibility of observing distributions $G_i^I(m)$ and $G_i^{II}(m)$ from the ideal scenario) which is, as discussed above, the key to identification. Prior to 2004, there were three PPO options under the old benefits and five entirely different PPO options under the new benefits. Hence, workers were forced to make an active choice.

In practice, to estimate the model the authors impose a parametric structure on the distribution function $F_\lambda(\cdot)$ and on the form of unobserved heterogeneity. Because of the model's fairly complex structure, estimation is carried out via MCMC. The estimation exercise yields an average coefficient of absolute risk aversion of $1.9 \cdot 10^{-3}$, with a large standard deviation of $2 \cdot 10^{-3}$. However, in this exercise risk aversion is no longer specified over financial outcomes, but over realized utility. As such, these estimates are not directly comparable to those obtained in the papers discussed earlier in this section.

The average value for ω is about \$1,300, which corresponds to one-third of the average health risk (about \$4,340) per employee-year. In other words, on average, moving from no insurance to full insurance yields an increase in medical utilization of \$1,300.

The authors interpret the parameter ω as (a measure of) moral hazard. In essence, ω maps the change in health insurance coverage into a change in utilization rate. And just like in the textbook version of adverse selection, *ceteris paribus*, individuals with higher moral hazard (higher ω) will have higher willingness to pay for insurance and will be more costly to insure.

6. Estimation with Aggregate Data

In this section, we turn our attention to research that estimates risk preferences using aggregate data. As in the prior section, we first discuss the general approach, and then review how that approach has been applied in several different contexts. Aggregate data has the major advantage that it is more easily obtained than individual-level data. As will become clear, however, estimation of risk preferences using aggregate data typically requires stronger assumptions about the underlying environment.

6.1 The General Approach

Estimating risk preferences from aggregate data requires steps that are similar to the ones faced when working with individual-level data, though with more structure imposed:

Step 1: Identify a field context in which economic agents make choices between options that involve risk, and where those choices determine some observed market aggregate for which the researcher can obtain data. Examples include asset prices, consumption aggregates, and labor supply elasticities.

Step 2: Translate the (typically) rich field environment into a well-defined market equilibrium model. Specifically, develop a model of choice (as in step 2 for individual-level

data) and a model of market forces such that, for the model of risk preferences under consideration, the market equilibrium model generates a prediction for the observed equilibrium market aggregates as a function of the taste parameters.

When using aggregate data, researchers frequently assume a representative agent, sometimes for simplicity—in particular, to make it easier to derive the mapping from model parameters to observed market aggregates—and sometimes due to there being limited information in the data from which one could identify either observed or unobserved heterogeneity. Recently, however, some researchers have shown that if one observes (arguably) the same set of agents interacting in multiple markets, then one can leverage techniques from the empirical industrial organization literature to estimate a model with heterogeneity in beliefs or in preferences (see section 6.2.2).

6.2 Betting Markets

A large literature estimates risk preferences using aggregate odds data from pari-mutuel horse races. In a pari-mutuel horse race, the total amount wagered net of a house take is distributed among the winning bettors in proportion to their individual bets. Specifically, suppose there are n horses in a race, and the share of the betting pool wagered on horse i is s_i . If the house take is τ , the net return on a one-dollar bet on horse i , or equivalently the “odds” on horse i , is given by

$$(8) \quad R_i = \frac{1-\tau}{s_i} - 1.$$

Typically, researchers working in this field context have access to data that contain, for each race, an observed vector of odds $\mathbf{R} = (R_1, \dots, R_n)$ and the identity of the

winning horse (and sometimes also the identity of the other horse placements). The literature takes advantage of the fact that, from \mathbf{R} , one can mechanically compute the house take and the vector of betting shares. Specifically, using the fact that the betting shares must sum to one,

$$(9) \quad 1 - \tau = \left(\sum_{i=1}^n \frac{1}{1 + R_i} \right)^{-1}.$$

The betting shares are then computed using equation (8):

$$(10) \quad s_i = \frac{(1 - \tau)}{(1 + R_i)}.$$

Hence, the literature typically takes \mathbf{R} , \mathbf{s} , and τ to be observed.⁵⁵

The conceptual framework behind the estimation of preferences from such data is as follows. The underlying primitives of a race are the house take τ and a vector of win probabilities $\mathbf{p} \equiv (p_1, \dots, p_n) \gg 0$, where p_i is the probability that horse i wins the race.⁵⁶ Given a specific vector of odds \mathbf{R} , bettors' preferences determine which horse they bet on, which in turn determines a set of market shares $\mathbf{s} \equiv (s_1, \dots, s_n)$. In equilibrium, the odds must adjust such that the shares implied by bettor behavior correspond to the shares implied by the pari-mutuel pricing rule (equations (9) and (10)). The goal is to use the observed equilibrium odds and race outcomes to make inferences about the unobserved preferences.

⁵⁵To be precise, while a specific \mathbf{s} can be associated with multiple \mathbf{R} via different τ , any observed \mathbf{R} is associated with a unique pair (\mathbf{s}, τ) .

⁵⁶Here in section 6.2, in the interest of using notation consistent with that in the betting-markets literature, we make a temporary shift in our notation and use p to denote probabilities (as opposed to prices).

6.2.1 Representative Agent Framework

Most studies that use pari-mutuel horse race data adopt a representative agent framework. The key implication is that, for each race, the equilibrium odds \mathbf{R} must be such that the representative agent is indifferent between betting on all horses in the race.

The earliest research in this area effectively tests whether bettors maximize expected value. In a seminal paper, Griffith (1949) uses data on a sample of 1,386 US horse races to study bettors' subjective risk perceptions. Griffith takes bettors' subjective win probabilities (which he labels the "psychological odds") to be the reciprocal of the track odds corrected for the house take.⁵⁷ He then compares these inferred subjective win probabilities to objective win probabilities. To do so, he groups horses into eleven odds-groups (e.g., all horses that went off at $R_i = 1$, $R_i = 2$, etc.), and he compares in each group the empirical win frequencies and the subjective win probabilities. He finds that the subjective win probabilities exceed the frequency of winners at short odds (corrected odds less than 5.1-to-1), while the reverse holds at long odds. This phenomenon has come to be known as the "favorite-longshot bias."⁵⁸

Weitzman (1965) takes a somewhat different approach, effectively using horse race data to estimate an EU model. Like Griffith (1949), he first estimates the relationship between empirical win frequencies and a

⁵⁷In particular, if the subjective win probability is taken to be $(1 - \tau)/(R_i + 1)$, then the expected value is the same (and equal to $1 - \tau$) for bets on all horses. Griffith (1949) also adds a correction for "breakage," wherein tracks round the odds downward to the nearest \$0.05.

⁵⁸Griffith (1961) reports similar findings in a follow-up study in which he turns attention from win bets to show bets. McGlothlin (1956) pursues a similar approach to Griffith (1949) and reaches the same conclusion using data from 9,248 California horse races. He further identifies changes in behavior over the course of a day's races that he takes to be inconsistent with the EU model.

horse's odds, using data from over 12,000 New York horse races. Specifically, for all 257 observed values of odds R , he computes the empirical frequency of winners p .⁵⁹ He then smooths this relationship by taking the 257 (R, p) pairs and fitting them to a curve $\hat{p}(R)$ using weighted least squares.⁶⁰ Unlike Griffith (1949), who then uses the odds to infer the representative agent's beliefs, Weitzman assumes that the representative agent has beliefs that correspond to $\hat{p}(R)$. Upon assuming that the agent behaves according to the EU model, he can infer the agent's utility function over money. In particular, normalizing $u(0) = 0$, indifference between all bets implies $\hat{p}(R)u(R) = C$, or $u(R) = C/\hat{p}(R)$, where C is a constant (reflecting the scale of utility). The estimated shape of $\hat{p}(R)$ implies that $u(R)$ is convex (increasing marginal utility for money), suggesting a region of local risk loving similar to that proposed by Markowitz (1952).

The first wave of papers culminates with Ali (1977). Using data on more than 20,000 New York harness races, Ali repeats the analyses of both Griffith (1949) and Weitzman (1965). However, he takes a different approach to grouping horses. In particular, he points out that the approach to estimating empirical win probabilities as a function of odds R ignores the fact that the relationship might depend on which other horses are in the race—e.g., a horse with 3-to-1 odds might be the favorite, the second favorite, or perhaps even the third favorite, and might thus have different win probabilities in each of those cases. To account for this, he instead groups horses according to “favorites,” where each race will have a favorite, second favorite, third favorite, and so

forth.⁶¹ It is unclear whether this alternative grouping is an improvement and, more to the point, it is unclear what underlying model of preferences might make this particular grouping the relevant one. However, the point Ali raises about the race composition is important, and it is picked up by Jullien and Salanié (2000).

As in Griffith (1949), Ali first estimates the subjective and empirical win probabilities for each of his groupings. He again finds evidence of a favorite–longshot bias. Specifically, the subjective win probability for the favorite is less than the objective win probability, whereas for all other horses the subjective win probability is larger than the objective win probability. As in Weitzman (1965), Ali reverts to assuming that the representative agent's beliefs correspond to the objective win probabilities. To generate the odds associated with each group, he computes a weighted average of the actual odds for that group. He can then infer $u(R)$ for each of these values of R . As in Weitzman (1965), the estimated shape of $u(R)$ is convex. By pursuing both approaches, the paper foreshadows the identification problem of inferring preferences versus beliefs.

In addition to reporting these empirical findings, Ali makes two important theoretical points. First, he observes that the utility function of the representative bettor is inconsistent with the EU hypothesis when the betting opportunity set is expanded to include parlay, martingale, and other compound bets (a point that is leveraged by Snowberg and Wolfers 2010; see below). Second, he proves that the favorite–longshot bias can be explained with bettors that are risk neutral EU maximizers but have heterogeneous risk perceptions (a point that is

⁵⁹In his data, odds are observed in \$0.05 increments.

⁶⁰He tests several functional forms and concludes that the most appropriate function is the hyperbolic form, $\hat{p}(R) = A/R + B \log(1 + R)/R$.

⁶¹Because the vast majority of races in his sample have, at most, eight horses, Ali (1977) carries out his analysis for the first eight favorites.

developed by Gandhi and Serrano-Padial 2015; see below).

Jullien and Salanié (2000) initiated a second wave of research using pari-mutuel horse race data that uses more sophisticated econometric techniques to estimate risk preferences. They use data from more than 34,000 horse races in Great Britain. In the spirit of Weitzman (1965), they assume that bettors' beliefs correspond to the objective win probabilities, and then estimate several different models of risk preferences, namely, EU, RDEU, and CPT.

Jullien and Salanié had a key innovation that solves the issue raised by Ali (1977) that the probability of a horse winning a race depends on the other horses in the race. Specifically, they use the fact that, as long as bettors' preferences satisfy certain conditions, any vector of observed odds $\mathbf{R} \equiv (R_1, \dots, R_n)$ implies a unique underlying vector of win probabilities $\mathbf{p} \equiv (p_1, \dots, p_n)$. Hence, if one can derive the mapping from \mathbf{R} to \mathbf{p} , one can use maximum likelihood to find the preferences that make the model-implied win probabilities as close as possible to the empirical win frequencies observed in the data on a race-by-race basis.

Formally, they assume a large number of identical bettors. Each bettor bets a fixed amount a in each race (and doesn't consider the possibility of not betting), bets that entire amount on one horse, and these bets cover every horse in the race.⁶² The bettors have identical preferences wherein the utility from betting a on horse i is $V(R_i, p_i, \theta)$, where θ is a vector that parametrizes preferences—i.e., for any particular model of risk preferences that they consider, the goal will be to estimate the vector θ . Under these assumptions, the market odds \mathbf{R} in each race

will adjust such that the utility from betting on all horses in that race is the same:

$$(11) \quad V(R_1, p_1, \theta) = \dots = V(R_n, p_n, \theta).$$

Jullien and Salanié argue that, as long as V is continuous and increasing in R_i and p_i (i.e., satisfies first-order stochastic dominance), then each pair (\mathbf{R}, θ) generates a unique vector of win probabilities \mathbf{p} .⁶³

With this result in hand, Jullien and Salanié (2000) can estimate θ using maximum likelihood. For each race, they number horses so that horse 1 is the winner. Indexing races by $c = 1, \dots, C$, the log-likelihood for the sample is

$$L^C(\theta) = \sum_{c=1}^C \log p_1(\mathbf{R}^c, \theta),$$

where \mathbf{R}^c are the observed odds for race c and $p_1(\mathbf{R}^c, \theta)$ is the objective win probability for the winning horse implied by \mathbf{R}^c and θ .

Jullien and Salanié (2000) first apply this approach to EU models. Here, the market equilibrium condition (11) implies

$$\begin{aligned} p_i u(w + aR_i, \theta) + (1 - p_i)u(w - a, \theta) \\ = \text{constant} \quad \text{for all } i = 1, \dots, n, \end{aligned}$$

where w is the bettor's initial wealth. Using these n equations and the fact that $\sum_{i=1}^n p_i = 1$, they obtain

$$p_i(\mathbf{R}, \theta) = \frac{[u(w + aR_i, \theta) - u(w - a, \theta)]^{-1}}{\sum_{j=1}^n [u(w + aR_j, \theta) - u(w - a, \theta)]^{-1}}.$$

⁶³Jullien and Salanié (2000) formally prove that, for any θ , there is a one-to-one relationship between \mathbf{R} and \mathbf{p} . However, they do not explicitly address the conditions required for how V depends on θ to ensure a one-to-one relationship between (\mathbf{R}, θ) and \mathbf{p} , which is what would be needed for full identification of θ . For models in which θ is unidimensional and changes in θ have a monotonic impact on risk aversion—as in single parameter EU models—the latter would clearly be satisfied. For models in which θ is multidimensional, it might not be satisfied.

⁶²Jullien and Salanié (2000) emphasize that their approach, in principle, allows for insider traders and noise traders, as long as the share of such traders is small enough that they don't influence the market odds.

Estimation of the EU model with a CARA utility function $u(x, \theta) = (1 - \exp(-\theta x))/\theta$ (for which only θa is identified) yields an estimate of $\hat{\theta}a = -0.055$, thereby indicating that bettors are risk loving. This result is consistent with the findings in Weitzman (1965) and Ali (1977), and follows from the fact that their data again contains evidence of a favorite–longshot bias. Similar results are obtained with a HARA utility function.

Jullien and Salanié (2000) next estimate various RDEU models. With RDEU, the market equilibrium condition implies

$$\begin{aligned} \pi(p_i, \theta)u(w + aR_i, \theta) \\ + (1 - \pi(p_i, \theta))u(w - a, \theta) = \text{constant} \end{aligned}$$

for all $i = 1, \dots, n$.

Much as under EU, one can use these n equations, the inverse function of $\pi(\cdot, \theta)$, and the fact that $\sum_{i=1}^n p_i = 1$ to obtain $p_i(\mathbf{R}, \theta)$. They assume a CARA utility function, and functional forms for $\pi(\cdot, \theta)$ include a power function and those proposed by Lattimore, Baker, and Witte (1992) and Prelec (1998).⁶⁴ The estimation results indicate that a power function and the Lattimore, Baker, and Witte (1992) function do not fit the data better than the EU model (in the sense of value of the likelihood function), and that the estimate $\hat{\theta}a$ is essentially unaffected. Estimation of the Prelec (1998) model yields a better fit to the data and suggests rejection of EU in favor of RDEU. However, the estimated weighting function is quite close to the diagonal.

Finally, Jullien and Salanié (2000) estimate CPT models. With CPT, and assuming that

⁶⁴In fact, they separately estimate the Lattimore, Baker, and Witte (1992) function and the Cicchetti and Dubin (1994) function, but the latter is equivalent to the former (with a simple change in variables).

w is the reference point, the market equilibrium condition implies

$$\begin{aligned} \pi^+(p_i, \theta)u(aR_i, \theta) + \pi^-(1 - p_i, \theta)u(-a, \theta) \\ = \text{constant} \quad \text{for all } i = 1, \dots, n. \end{aligned}$$

The approach to solving for $p_i(\mathbf{R}, \theta)$ is analogous, except that permitting π^+ to differ from π^- requires that they solve for $p_i(\mathbf{R}, \theta)$ numerically rather than analytically.⁶⁵ They again assume a CARA utility function, and functional forms for $\pi^+(\cdot, \theta)$ and $\pi^-(\cdot, \theta)$ include a power function and that proposed by Lattimore, Baker, and Witte (1992).⁶⁶ In each case, the estimate $\hat{\theta}a$ is not far from what is obtained for the EU model, the probability weighting function for gains is slightly convex but not significantly so, but the probability weighting function for losses is highly and significantly concave, leading to a clear rejection of EU.

Snowberg and Wolfers (2010) revisit the favorite–longshot framing of the data as in Griffith (1949) and Weitzman (1965), and investigate whether it is driven by risk love (increasing marginal utility) or by risk misperceptions (probability distortions). As in Ali (1977), either model can fully explain data on win bets. The fundamental insight used for identification is that the so-estimated models will have different implications for bettors' decisions over a wider choice set (as suggested by Ali 1977), in particular for compound bets in the exacta, quinella, or trifecta pools.⁶⁷ Snowberg and Wolfers's data cover

⁶⁵Given computing power at the time, they thus used a small subset of the data and considered a more limited set of models.

⁶⁶By assuming a CARA utility function, Jullien and Salanié (2000) effectively assume that there is no diminishing sensitivity or loss aversion (see section 3.3). In particular, the estimates would be the same if $u(aR_i, \theta)$ and $u(-a, \theta)$ were replaced with $u(w + aR_i, \theta)$ and $u(w - a, \theta)$.

⁶⁷An exacta is a bet on both which horse will come in first and which will come in second. A quinella is a bet on two horses to come in first and second in either order. A

every horse race run in North America from 1992 to 2001, including the finishing position of each horse, the win odds on each horse, and the realized payoffs for exactas, quinellas, and trifectas. As in the prior literature, Snowberg and Wolfers take a representative-agent approach: they assume a large number of identical bettors, where each bets a fixed amount of \$1 in each race (or can choose not to bet), and bets that \$1 on (at most) one horse. Hence, again the market odds will be such that the utility of all bets (and the utility of not betting) are all the same.

The analysis proceeds in three steps. First, Snowberg and Wolfers (2010) use race results to estimate how the win probability depends on the odds. They pursue an approach similar to Weitzman (1965), first calculating the empirical frequency of winners for each value of odds R observed in the data, and then using Lowess smoothing to estimate a function $\hat{p}(R)$.⁶⁸

Next, Snowberg and Wolfers (2010) separately fit the EU model and the risk misperception model to the estimated $\hat{p}(R)$ function. For the EU model, they assume that bettors' beliefs correspond to $\hat{p}(R)$, in which case, the market equilibrium condition implies

$$\begin{aligned} \hat{p}(R_i)u(w + R_i) + (1 - \hat{p}(R_i))u(w - 1) \\ = u(w) \quad \text{for all } i = 1, \dots, n. \end{aligned}$$

Normalizing $w = 0$, normalizing utility to zero if the bet is lost ($u(-1) = 0$),

trifecta is a bet on which horse will come in first, which second, and which third.

⁶⁸They actually estimate a smooth function for the rate of return as a function of R , but since the rate of return is equal to pR , this is equivalent to estimating a smooth function for the win probability $\hat{p}(R)$.

and normalizing utility to one if the bettor chooses not to bet ($u(0) = 1$), this becomes

$$(12) \quad \hat{p}(R_i)u(R_i) = 1 \Rightarrow u(R_i) = \frac{1}{\hat{p}(R_i)},$$

$$i = 1, \dots, n.$$

In other words, a utility function $u(R) = 1/\hat{p}(R)$ perfectly fits the estimated $\hat{p}(R)$ function. The so-estimated $u(R)$ is, not surprisingly, convex—as we have seen, explaining the favorite–longshot bias with an EU model requires a risk-loving utility function. While they prefer (and use below) this non-parametric estimate for $u(R)$, they also argue that a CARA utility function fits the data reasonably well.

For the risk-misperception model, they assume that bettors have a linear utility function $u(x) = x$. Letting $\pi(p)$ denote the misperceived win probability associated with objective win probability p , the market equilibrium condition implies

$$\pi(\hat{p}(R_i))R_i + (1 - \pi(\hat{p}(R_i)))(-1) = 0$$

$$\text{or } \pi(\hat{p}(R_i))(R_i + 1) = 1$$

$$\text{for all } i = 1, \dots, n.$$

Letting $\hat{R}(p)$ denote the inverse of $\hat{p}(R)$, it follows that a risk-misperception function $\pi(p) = 1/(\hat{R}(p) + 1)$ perfectly fits the estimated $\hat{p}(R)$ function. The so-estimated $\pi(p)$ indicates significant overestimates of small probabilities, exactly what is needed to explain the favorite–longshot bias. Again, while they prefer (and use below) this non-parametric estimate for $\pi(p)$, they also argue that a one-parameter probability weighting function as in Prelec (1998) fits the data reasonably well.

Finally, in step 3, having established that either model can perfectly fit data on win bets, Snowberg and Wolfers (2010) turn to compound bets. They (and we) outline their approach only for exacta bets (in which the better wagers that horse A will come in first and horse B will come in second); analogous approaches are used for quinella and trifecta bets. For exacta bets, they do not observe the odds for all combinations of horses, but rather only the odds for the winning combination. Hence, they use the two estimated models from step 2 to generate, for each race, model-predicted exacta odds \hat{E}_{AB} for the winning combination, and then compare those to the observed exacta odds E_{AB} . In other words, they investigate which estimated model from step 2 best predicts the observed exacta odds.

Letting $p_{B|A}$ denote the probability that horse B is second given that horse A wins, the predicted exacta odds can be derived much as in step 2:⁶⁹

EU:

$$p_A p_{B|A} u(\hat{E}_{AB}) = 1 \Rightarrow \hat{E}_{AB} = u^{-1}\left(\frac{1}{p_A p_{B|A}}\right);$$

Risk Misperception:

$$\begin{aligned} \pi(p_A)\pi(p_{B|A})(\hat{E}_{AB} + 1) &= 1 \\ \Rightarrow \hat{E}_{AB} &= \frac{1}{\pi(p_A)\pi(p_{B|A})} - 1. \end{aligned}$$

For each race, p_A and $\pi(p_A)$ can be derived from R_A as in step 2. The key question is how to derive $p_{B|A}$ and $\pi(p_{B|A})$. Snowberg and Wolfers pursue two approaches, both of which yield the same conclusions. First, they use Harville's (1973) formula,

which yields $p_{B|A} = p_B / (1 - p_A)$ or $\pi(p_{B|A}) = \pi(p_B) / (1 - \pi(p_A))$, where p_B and $\pi(p_B)$ can be derived from R_B as in step 2.⁷⁰ Second, they estimate $\hat{p}(R_B|R_A)$ using a technique analogous to how they estimate $\hat{p}(R)$.

They compare models by computing the mean absolute prediction error between \hat{E}_{AB} and E_{AB} for each model. Using this metric, they find that the estimated risk-misperception model better predicts observed odds for exacta bets (and for quinella and trifecta bets). The authors thus conclude that their "results are more consistent with the favorite-long shot bias being driven by misperceptions rather than by risk love" (p. 744).

6.2.2 Heterogeneity in Beliefs and Preferences

Two recent papers—Gandhi and Serrano-Padial (2015) and Chiappori et al. (2012)—investigate the extent to which one can use aggregate data on horse races to estimate heterogeneity in beliefs and heterogeneity in preferences. The key idea in both is that, instead of assuming a representative bettor who is indifferent between betting on all horses in a race, it is assumed that heterogeneous bettors partition themselves across the horses in a race based on either their beliefs or their preferences. If one can observe how bettors partition themselves across races with different characteristics, one can draw inference on the distribution of heterogeneity.

Gandhi and Serrano-Padial (2015) are motivated by the favorite-longshot bias, and the idea that belief heterogeneity could drive it (as suggested by Ali 1977, theorem 2). They develop a formal model of risk-neutral agents with heterogeneous beliefs, and prove that such a model would generate a favorite-longshot bias. Here, we focus on the

⁶⁹The risk-misperception model here assumes that bettors fail to reduce the compound lottery inherent in an exacta—i.e., they do not use $\pi(p_A p_{B|A})$.

⁷⁰Under Harville's model, the ratio of the probabilities that any two horses place second does not depend on which of the other horses wins.

latter part of the paper, where they estimate the extent of belief heterogeneity.

Once again, a race consists of a house take τ and an objective vector of win probabilities $\mathbf{p} \equiv (p_1, \dots, p_n)$, and all bettors bet the same fixed amount. Bettor t holds beliefs $\varphi_t \equiv (\varphi_{1t}, \dots, \varphi_{nt})$, where beliefs are heterogeneous and thus some bettors have $\varphi_t \neq \mathbf{p}$. Letting $Q(\varphi_t; \theta)$ denote a population distribution of φ_t , the goal is to estimate the vector of parameters θ that characterize Q . Since bettors are risk neutral, bettor t bets on the horse with the largest $\varphi_{it}R_i$. Assuming that Q is continuous, so that a measure zero of bettors are indifferent between multiple horses, any vector of odds \mathbf{R} will cause bettors to partition themselves such that the share betting on horse i is

$$s_i = \Pr_Q(\varphi_{it}R_i > \varphi_{jt}R_j \text{ for all } j \neq i).$$

In equilibrium, odds must adjust such that the shares implied by this bettor behavior correspond to the shares implied by the pari-mutuel pricing rule (equations (9) and (10)).

Gandhi and Serrano-Padial (2015) had the key insight that this market is isomorphic to a discrete-choice horizontally differentiated products market of the form studied by Berry, Levinsohn, and Pakes (1995) and Berry, Gandhi, and Haile (2013).⁷¹ Specifically, $\varphi_{it}R_i > \varphi_{jt}R_j$ is equivalent to $\delta_i + \nu_{it} > \delta_j + \nu_{jt}$ where $\delta_i = \log p_i R_i - \log p_1 R_1$ and $\nu_{it} = \log(p_1 / \varphi_{1t}) - \log(p_i / \varphi_{it})$. Hence, instead of focusing on the population distribution $Q(\varphi_t; \theta)$, they focus on the population distribution $P(\boldsymbol{\nu}_t; \theta)$. The market shares can now be written as

$$s_i = \Pr_P(\delta_i + \nu_{it} > \delta_j + \nu_{jt} \text{ for all } j \neq i),$$

⁷¹In an appendix, they describe how this approach can be extended to permit risk-averse bettors.

which is the familiar form in the discrete-choice literature. Applying results from that literature, for any continuous $P(\boldsymbol{\nu}_t; \theta)$, one can translate any observed odds \mathbf{R} into a unique vector of model-implied underlying probabilities $\mathbf{p}^*(\mathbf{R}, \theta)$. One can then set up a likelihood function exactly as in Jullien and Salanié (2000).

The authors estimate this model—i.e., they estimate $P(\boldsymbol{\nu}_t; \theta)$ —using a sample of more than 176,000 pari-mutuel races that were collected from North American tracks over 2003–06, consisting of more than 1,400,000 horse starts.⁷² Much as in the discrete-choice literature, they assume that $P(\boldsymbol{\nu}_t; \theta)$ takes a logistic form. However, they permit a variance mixture of logistic functions. Specifically, they assume

$$P(\boldsymbol{\nu}_t; \theta) = \sum_{m=1}^M \left(\prod_{i=2}^n F(\nu_{it} | \sigma_m) \right) g(\sigma_m),$$

where F is a standard logistic distribution with variance σ_m , and $g(\sigma_m)$ is the probability that the population variance equals σ_m .⁷³ They interpret σ_m to be a measure of belief heterogeneity for type m . The authors report results for $M = 1, 2$, and 3. A likelihood ratio test rejects the one-type model in favor of the two-type model, while adding a third type does not significantly improve the log-likelihood. In the two-type model, they estimate that 70 percent of the population have a small $\sigma = 0.028$, while 30 percent have a large $\sigma = 0.503$. The authors interpret the former group as having roughly correct beliefs (informed traders), and they interpret the latter group as having dispersed beliefs (noise traders).⁷⁴

⁷²As a preliminary step, the authors regress *ex post* return on (log) market shares, and indeed find evidence of the favorite-longshot bias in their data.

⁷³The authors don't discuss how this functional-form assumption translates into a distribution of the underlying beliefs φ_t .

⁷⁴They also estimate a model with risk-averse bettors who all have the same CARA utility function, and reach

Finally, the authors do some additional analysis to further buttress their argument that belief heterogeneity is playing an important role in this context. First, they use their data to estimate a representative agent model using the preferred specification of Jullien and Salanié (2000), and they conclude using a Vuong test that their belief-heterogeneity model is statistically better than the representative agent model. Second, they separately estimate their model on the subsamples of maiden and non-maiden races. Maiden races are races in which participating horses have yet to win a single race, and thus there is less handicapping information about horses in maiden races than in non-maiden races. Their estimates suggest stable proportions of informed versus noise traders across the two types of races, but lower variance estimates for each type of trader in the non-maiden races. This result is exactly what one might expect if belief heterogeneity were playing an important role in this environment.

Chiappori et al. (2012) take the opposite approach to heterogeneity: they assume that bettors have correct (and thus homogeneous) beliefs, and estimate a model with heterogeneity in preferences. Yet again, a race consists of a house take τ and an objective vector of win probabilities \mathbf{p} , where the authors label the horses in each race such that $p_1 > \dots > p_n$. They assume a continuum of bettors who all have beliefs that correspond to \mathbf{p} and who all bet the same fixed amount. Bettors differ in their utility function, $V(R, p, \theta)$, which is the utility from betting on a horse that has odds R and win probability p . The parameter θ captures heterogeneity.

The authors do not put any parametric structure on V , but rather assume that it satisfies a few simple properties. They assume that V is continuous and increasing in R and

p . They assume that θ is unidimensional (i.e., θ is a scalar), and without loss of generality that it is distributed uniformly on the interval $[0, 1]$. Most importantly, V satisfies a single-crossing property with respect to θ : for any two gambles (R, p) and (R', p') with $p' < p$, if, for some θ , $V(R, p, \theta) \leq V(R', p', \theta)$, then for all $\theta' > \theta$, $V(R, p, \theta') < V(R', p', \theta')$. Hence, θ can be interpreted as a taste for risk, where the larger θ is, the more a bettor is prone to prefer horses with a lower probability but higher odds.⁷⁵

Given this structure, for any fixed V , \mathbf{p} , and \mathbf{R} such that money is bet on every horse (which must hold in equilibrium), bettors partition themselves by θ : there exists $0 = \theta_0 < \theta_1 < \dots < \theta_n = 1$ such that bettors with $\theta \in (\theta_0, \theta_1)$ bet on horse 1, bettors with $\theta \in (\theta_1, \theta_2)$ bet on horse 2, and so on until bettors with $\theta \in (\theta_{n-1}, \theta_n)$ bet on horse n . In other words, bettors with the lowest θ s, who are most averse to risk, bet on the horse with the largest win probabilities, and bettors with larger and larger θ s bet on horses with smaller and smaller win probabilities. It follows that market shares as a function of odds \mathbf{R} satisfy $s_i = \theta_i - \theta_{i-1}$. In equilibrium, odds must adjust such that the shares implied by this bettor behavior equal the shares implied by the pari-mutuel pricing rule (equations (9) and (10)).

The authors prove that, for any V satisfying their assumptions, for each race (\mathbf{p}, τ) there exists a unique set of equilibrium odds \mathbf{R} , and moreover for any observed equilibrium odds \mathbf{R} , there is a unique vector of win probabilities \mathbf{p} that could generate those odds. Finally, the authors prove that, if one observes the relationship $\mathbf{p}(\mathbf{R})$, one can identify the utility function V up to an affine

much the same conclusion about the nature of belief heterogeneity.

⁷⁵This single-crossing assumption is satisfied for an EU model with CARA utility or with CRRA utility. For non-EU models, it can be satisfied, though additional restrictions are required. In particular, heterogeneity has to remain unidimensional.

transformation. Hence, in principle one can nonparametrically identify V , and the extent of heterogeneity via V 's dependence on θ , from the aggregate pricing data.

To relate the nonparametric utility function V to various underlying models of preferences, the authors focus on a feature of V known as the “normalized fear of ruin”:

$$NF(R, p, \theta) = \frac{p}{R+1} \frac{V_p}{V_R}(R, p, \theta).$$

Different models of risk preferences have different implications for $NF(R, p, \theta)$ —e.g., under EU it is independent of p . To estimate NF , they propose a two-step procedure. First, they use empirically observed outcomes to nonparametrically estimate the function $p_i(\mathbf{R})$.⁷⁶ Second, with the function $p_i(\mathbf{R})$ in hand, they are able to estimate the NF index at each observation in the data.⁷⁷ They implement this approach using a data set of more than 53,000 thoroughbred races in the United States from 2001 to 2004. They then test how well different models of risk preferences can explain the variation in the estimated NF index. Of the models considered, a heterogenous non-EU model with non-additive probability weights, in which the estimated probability weighting function is concave and then convex, performs best.

The approach in Chiappori et al. (2012) has limitations. In particular, the assumption of unidimensional heterogeneity that

⁷⁶This step is analogous to the approach in Weitzman (1965) and Snowberg and Wolfers (2010), except that they estimate how the probability of horse i winning depends on the full vector of odds \mathbf{R} , as opposed to merely depending on R_i .

⁷⁷Here, they do not pursue a maximum-likelihood approach as in Jullien and Salanié (2000) and Gandhi and Serrano-Padial (2015). Rather, they focus on one particular implication of V that can be non-parametrically identified from $\mathbf{p}(\mathbf{R})$, and which can be used to compute NF . For details, see the paper.

satisfies the single-crossing property is quite restrictive for non-EU models. Also, their estimation approach is not as straightforward to implement as other approaches in this literature. That said, their intuition for how to identify heterogeneity in preferences from aggregate data is quite appealing.

6.3 Consumption, Asset Returns, and Labor Supply Data

A number of studies use macroeconomic data to estimate risk preferences. Such studies typically use a representative-agent life-cycle model, and use aggregate data to estimate a parameter of that model that reflects the representative agent's degree of risk aversion. Here, we describe a few seminal papers in this literature, focusing on papers in which a major focus is the estimation of a risk parameter. As we highlight below, however, the parameter estimated differs across papers.

Hansen and Singleton (1983) use data on consumption and asset returns to estimate a life-cycle consumption model. They study a single-good economy of identical, infinitely lived agents with time-additive EU preferences. The representative agent in this economy chooses a stochastic consumption plan $\{c_t\}_{t=0}^{\infty}$ to maximize

$$(13) \quad E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right],$$

where E_0 is the expectation conditional on information available in period zero, $\beta > 0$ is a discount factor, and $u(c_t)$ is the per-period utility function (which is monthly in the estimation). The representative agent faces the sequence of budget constraints

$$(14) \quad c_t + \mathbf{p}'_t \mathbf{x}_{t+1} \leq (\mathbf{p}_t + \mathbf{d}_t)' \mathbf{x}_t,$$

where \mathbf{x}_t is a vector of the agent's holdings of n assets in period t , \mathbf{p}_t is the vector of prices

of the n assets net of any distributions, and \mathbf{d}_t is the vector of distributions in period t .⁷⁸

Hansen and Singleton (1983) assume a CRRA utility function $u(c_t) = (c_t^{1-\rho} - 1)/(1 - \rho)$, and they estimate the coefficient of relative risk aversion ρ . In other words, they estimate utility curvature over per-period consumption. Given this functional form, the first-order necessary conditions for the maximization of (13) subject to (14) are given by the Euler equations,

$$(15) \quad E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} R_{it+1} \right] = 1, \quad i = 1, \dots, n,$$

where E_t is the expectation conditional on information available in period t and $R_{it+1} = (p_{it+1} + d_{it+1})/p_{it}$ is the one-period return on asset i .

By assuming, *inter alia*, that the joint distribution of consumption and asset returns is log-normal, Hansen and Singleton (1983) are able to obtain ML estimates of ρ and β using monthly US data for the period February 1959 through December 1978.⁷⁹ In general, their estimates of ρ range between zero and two, and their estimates of the discount factor β are less than but close to unity. Perhaps more importantly, however, they perform various chi-square and likelihood-ratio tests that provide substantial evidence against the model. Of course, these tests are joint tests of the model's several restrictions, including the preference assumptions (identical agents, EU preferences, time additivity, CRRA utility, exponential discounting, etc.) and

the distributional assumption (jointly log-normally distributed consumption and asset returns), making it impossible to say which restrictions are being rejected. For this reason, the authors lay out a dual plan of pursuing both “models with more general specifications of preferences and distribution-free methods of estimation and inference” (Hansen and Singleton 1983, p. 264).

In Hansen and Singleton (1982), the authors progress their plan to pursue distribution-free methods of estimation and inference.⁸⁰ They develop a generalized instrumental variables procedure for estimating the parameters of model (13)–(14) and implement their procedure using the same monthly data on consumption and asset returns used in Hansen and Singleton (1983). In brief, they first use the Euler equations (15) to generate a set of population orthogonality conditions,

$$(16) \quad E \left[\begin{pmatrix} \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} R_{1t+1} - 1 \\ \vdots \\ \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} R_{nt+1} - 1 \end{pmatrix} \otimes \mathbf{z}_t \right] = \mathbf{0},$$

where E is unconditional expectation and the vector of instruments \mathbf{z}_t comprises lagged values of $(R_{1t+1}, \dots, R_{nt+1}, c_{t+1}/c_t)'$. They then use the sample analog of the orthogonality conditions to construct a generalized method of moments (GMM) estimator of the model parameters. The GMM estimates of ρ and β are similar to the ML estimates reported in Hansen and Singleton (1983). Chi-square tests again provide evidence against the model, which here may be interpreted as direct evidence against the model's preference assumptions.

⁷⁸In Hansen and Singleton (1983), the right-hand side of the budget constraint also includes a term explicitly measuring the agent's labor income in period t . This term can be suppressed, however, without loss of generality (Epstein and Zin 1989, 1991, p. 267).

⁷⁹For aggregate consumption, they use seasonally adjusted real consumption at the monthly level. They estimate the model using several different monthly asset return series.

⁸⁰Hansen and Singleton (1982) was published before, but apparently written after, Hansen and Singleton (1983).

The rejection by Hansen and Singleton (1982, 1983) of the standard-preference assumptions is echoed in the subsequent literature on the equity premium puzzle (Mehra and Prescott 1985).⁸¹ In their famous paper, Mehra and Prescott consider a variation of the standard model studied by Hansen and Singleton in which there are two assets: a risky equity security and a risk-free debt security. Unlike Hansen and Singleton, however, Mehra and Prescott do not estimate the model's preference parameters. Instead, they calibrate the model's non-preference parameters so that the population mean, variance, and autocorrelation of consumption growth match their corresponding sample values for the US economy between 1889 and 1978. They then derive analytical expressions for the expected returns on the equity and risk-free securities, R^e and R^f , in terms of the model's preference parameters, ρ and β . Restricting $\rho \leq 10$, $\beta \leq 1$, and $R^f \leq 0.04$ (more than four times the average return on the three-month Treasury bill during the sample period), they find that the maximum equity premium ($R^e - R^f$) that is consistent with the model is less than one percent. In contrast, in their data, the historical equity premium (the difference between the average annual returns on the S&P 500 and the three-month Treasury bill) was more than six percent.

Motivated in part by the poor empirical performance of the standard model in Hansen and Singleton (1982, 1983) and Mehra and Prescott (1985), and in part by concurrent developments in the microeconomics literature on non-EU preferences, Epstein and Zin (1989, 1991) pursue a model with a more general specification of preferences.⁸² The key innovation of their

model is that it disentangles the representative agent's degree of risk aversion from his elasticity of intertemporal substitution. In the standard model, the coefficient of relative risk aversion is constrained to equal the reciprocal of the elasticity of intertemporal substitution. As a result, an agent with standard preferences that is very averse to risk must also have a very strong preference for consumption smoothing. However, this is what makes the equity premium a puzzle: the high degree of risk aversion that is required to explain the historical equity premium implies an implausibly low elasticity of intertemporal substitution, given the historical rate of consumption growth. Accordingly, Epstein and Zin investigate whether a model that delinks preferences over risk and intertemporal substitution can outperform the standard model.

In Epstein and Zin's model, the representative agent's utility in period t is defined recursively as

$$(17) \quad U_t = \left[(1 - \beta) c_t^\varsigma + \beta (E_t [U_{t+1}]^{1-\rho})^{\frac{\varsigma}{1-\rho}} \right]^{\frac{1}{\varsigma}}$$

subject to the sequence of budget constraints given by (14). In this formulation, $\varsigma = 1 - (1/\sigma)$ where σ is the elasticity of intertemporal substitution, $\beta > 0$ is a discount factor, and ρ is the coefficient of relative risk aversion. Hence, as in Hansen and Singleton (1982, 1983), in terms of risk preferences Epstein and Zin (1991) estimate utility curvature over per-period consumption, though with a different functional form.

The Euler equations that characterize the first-order necessary conditions of the optimal consumption path are

$$(18) \quad E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{\theta(\varsigma-1)} (m_{t+1})^{\theta-1} R_{it+1} \right] = 1,$$

$$i = 1, \dots, n,$$

⁸¹For a survey of the literature on the equity premium puzzle, see Kocherlakota (1996). See also Siegel and Thaler (1997). Campbell (2003) summarizes the larger literature on asset price puzzles in macroeconomics.

⁸²Weil (1989, 1990) takes a parallel path.

where $\theta = (1 - \rho)/\zeta$ and $m_{t+1} = [(\mathbf{p}_{t+1} + \mathbf{d}_{t+1})' \mathbf{x}_{t+1}] / \mathbf{p}_t' \mathbf{x}_{t+1}$ is the one-period return on the agent's asset holdings in period t . Observe that when $\rho = 1/\sigma$, and thus $\zeta = 1 - \rho$, the Euler equations (18) reduce to (15).

Following the GMM procedure of Hansen and Singleton (1982), Epstein and Zin (1991) estimate the model using monthly US data for the period February 1959 through December 1986. They estimate the coefficient of relative risk aversion ρ to be close to one and the elasticity of intertemporal substitution σ to be less than one. Moreover, Wald and likelihood ratio tests of the restriction $\rho = 1/\sigma$ generally reject the standard model. That said, Epstein and Zin's (1989, 1991) model does not resolve the equity premium puzzle, in part reflected by the fact that the discount factor β is often estimated to be greater than one.⁸³

In a completely different approach, Chetty (2006) analyzes risk preferences using data on labor supply. Chetty's basic insight is that the wage elasticity of labor supply, which has been estimated in numerous studies in labor economics, provides information about the curvature of the marginal utility of consumption.

Chetty (2006) begins with a standard life-cycle labor-supply model in which a representative agent chooses a path of consumption $\{c_t\}_{t=1}^T$ and labor $\{l_t\}_{t=1}^T$ to maximize $U(c_1, \dots, c_T, l_1, \dots, l_T)$ subject to the budget constraint

$$\begin{aligned} p_1 c_1 + \dots + p_T c_T \\ = y + w(\alpha_1 l_1 + \dots + \alpha_T l_T), \end{aligned}$$

⁸³Another generalization of the standard EU model that features prominently in the macro-finance literature is the model with habit formation (e.g., Abel 1990, 1999; Constantinides 1990; Campbell and Cochrane 1999).

where p_t is the price of consumption in period t , y is unearned income at time 0, and $w\alpha_t$ is the wage in period t . Chetty chooses to work with the associated indirect utility function $u(c, l)$, where $c \equiv p_1 c_1 + \dots + p_T c_T$ represents lifetime consumption expenditures, $l \equiv \alpha_1 l_1 + \dots + \alpha_T l_T$ represents lifetime labor supply, and the budget constraint becomes $c = y + wl$. He then studies the coefficient of relative risk aversion in u — that is,

$$\gamma \equiv \frac{-u_{cc}(c, l)c}{u_c(c, l)}.$$

In other words, whereas Hansen and Singleton (1982, 1983) and Epstein and Zin (1989, 1991) study curvature of the per-period utility function over per-period consumption, Chetty studies curvature of the indirect utility function with respect to lifetime consumption (or, equivalently, lifetime wealth).

Given this model, the first order condition with respect to labor is given by

$$wu_c(y + wl, l) = -u_l(y + wl, l).$$

Chetty (2006) shows that, with some implicit differentiation and algebraic manipulation, one can derive that

$$\gamma = -\left(1 + \frac{wl}{y}\right) \frac{\varepsilon_{l,y}}{\varepsilon_{l,w}} + \left(1 + \frac{y}{wl}\right) \varepsilon_{u_c,l},$$

where $\varepsilon_{l,y}$ is the income elasticity of labor supply, $\varepsilon_{l,w}$ is the compensated wage elasticity of labor supply, and $\varepsilon_{u_c,l}$ is the elasticity of the marginal utility of consumption with respect to labor (which captures any complementarity between consumption and labor supply). Hence, if one has estimates of these three elasticities along with values for y and wl , one can obtain an estimate for γ .

In his empirical implementation, Chetty (2006) takes advantage of a large number of existing studies that contain estimates of the two needed wage elasticities along with measures of y and wl (in most cases he uses the sample means of y and wl). An estimate of the elasticity of the marginal utility of consumption with respect to labor is harder to come by, but he appeals to evidence on consumption responses to job loss, disability, and retirement to conclude that this elasticity is, at most, 0.15. Hence, for each existing study with estimates for the needed wage elasticities, Chetty provides the corresponding coefficient of relative risk aversion γ under an assumption of $\varepsilon_{u,c,l} = 0$ and of $\varepsilon_{u,c,l} = 0.15$. Overall, much as in Hansen and Singleton (1982, 1983) and Epstein and Zin (1989, 1991), he finds only modest risk aversion—the average estimate for γ is 0.71 in the former case and 0.97 in the latter case.

7. *Directions for Future Research*

In this section we discuss a number of directions for, and important questions to be addressed in, future research.

7.1 *Consistency across Contexts*

A common assumption in economics is that risk preferences are stable across decision contexts. Under EU, for instance, the implication is that multiple risky choices by the same agent should reflect the same degree of risk aversion (or risk loving). If this assumption is correct, then estimates of risk preferences derived from choices in one context can be used to understand and make predictions about the behavior of households in other contexts. Assessing the empirical validity of the stability assumption is a difficult task. Moreover, it is quite possible that risk preferences are stable across a certain set of contexts, but not others. While there exist a few papers on this

issue, which we survey below, more work is needed.

In an early paper, Wolf and Pohlman (1983) compare the risk preferences of a single person—a dealer in US government securities—in two contexts; a hypothetical gambling context and an actual investment context. For both cases, they assume that the dealer is an EU maximizer with HARA utility. In the gambling context, Wolf and Pohlman recover the utility parameters from the dealer's direct assessments of six hypothetical lotteries. In the investment context, they recover the same parameters from the dealer's bid decisions over a series of twenty-eight Treasury bill auctions conducted during a twenty-week interval in 1976 along with the dealer's own subjective forecasts of the distribution of short-term returns. They find that the dealer was about four times more risk averse in the actual bid decisions than he was for the hypothetical gambles, leading them to conclude that people's "degree of risk aversion may depend on the specific context in which their choices are made" (p. 849).

Barseghyan, Prince, and Teitelbaum (2011) investigate the stability hypothesis by examining the choices of 702 households across three insurance coverages: auto collision, auto comprehensive, and home all perils. As in BMOT, they observe for each choice the full menu of premium–deductible combinations, the household's chosen deductible, and the household's claim rate for that coverage (estimated from claims data and demographics). Households are assumed to have subjective beliefs that correspond to their coverage-specific claim rates and to be EU maximizers with an NTD utility function. The authors then focus on the stability within a household of the coefficient of absolute risk aversion r .

Specifically, each choice that a household makes implies an interval $[\underline{r}, \bar{r}]$ such that any $r \in [\underline{r}, \bar{r}]$ is consistent with that choice. For instance, if a household chooses a \$250

deductible on home insurance, there exists an \underline{r} such that any $r < \underline{r}$ would imply the household should choose a deductible larger than \$250, and there exists an \bar{r} such that any $r > \bar{r}$ would imply the household should choose a deductible smaller than \$250. The authors ask, for each household, whether the three intervals implied by its three choices intersect. If so, any r contained in the intersection can rationalize the household's choices (i.e., the household might have stable risk preferences); if not, no r can rationalize the household's choices.

Barseghyan, Prince, and Teitelbaum (2011) find that the hypothesis of stable risk preferences is rejected by the data. The three intervals intersect for only 23 percent of households, which is rather low considering that, according to the authors' calculations, 14 percent of households would have overlapping intervals even if households were randomly assigned their deductible choices. The authors pursue several more sophisticated approaches to demonstrate that the conclusion of limited stability is robust. The authors further point out that the major inconsistency seems to come from households tending to exhibit greater risk aversion in their home deductible choices than they do in their auto deductible choices.

Einav et al. (2012) also study the stability of risk preferences, although in contrast to Barseghyan, Prince, and Teitelbaum (2011), they focus on the null of no stability, and assess the extent to which there is a domain-general component to risk preferences. They use data on the benefits choices of 12,752 Alcoa employees in six contexts: health insurance, drug insurance, dental insurance, short-term disability insurance, long-term disability insurance, and 401(k) investments.

Einav et al. (2012) take two complementary approaches in their analysis. First, they take a model-free approach in which they rank by risk the options within each context and compute the pairwise rank correlations

in the employees' choices across contexts. They find that an employee's choice in every context is positively correlated with his or her choice in every other context, and thus conclude that they can reject the null hypothesis of no domain-general component of risk preferences. Second, they take a model-based approach that is conceptually similar to the approach in Barseghyan, Prince, and Teitelbaum (2011). They specify an EU model that permits both an individual-specific coefficient of risk aversion and a context-specific shift in risk preferences (that does not vary across employees). The latter effect is modeled in a stylized, reduced-form way that does not really correspond to any underlying model of risk preferences. They then choose the context-specific shifts that maximize the number of employees for whom there exists an individual-specific coefficient of risk aversion that, when combined with the context-specific shifts, can explain their behavior across all six contexts. With this approach, they can fully rationalize the behavior of roughly 30 percent of employees across all six contexts, which they take to be further evidence that there is a domain-general component of risk preferences (though they acknowledge that the size of the context-specific shifts "suggests that the implied levels of risk aversion exhibited may be very different across domains, or that other effects, such as framing or probability weighting, are particularly important in these contexts" (p. 2634)).

Barseghyan, Molinari, and Teitelbaum (2016), discussed in section 5.2, demonstrate a close connection between rank correlation of choices and stability of risk preferences under a probability distortion model. They find that stability of risk preferences cannot be rejected for roughly five in six households whose choices are rationalizable by the model. They then document that households' deductible choices are rank correlated across lines of coverage, echoing the finding

by Einav et al. (2012), and they show that it is stable households who are driving these rank correlations.

Moving forward, whether in an EU or a non-EU framework, it would be natural to consider models of risk preferences that have both a permanent (domain-general) and an idiosyncratic (context-specific) component. As a simple example, in an EU model with CRRA utility, one might permit that household i 's coefficient of risk aversion in choice c is given by $\rho_{ic} = \bar{\rho}_i + \varepsilon_{ic}$, where the $\bar{\rho}_i$ is constant across all of i 's choices, while ε_{ic} is drawn separately for each choice that household i faces. With appropriate data—and in particular with enough observations for each household—one could estimate the extent and importance of each.

7.2 Combining Experimental and Field Data

A special case of the question of consistency across contexts is the question of consistency between behavior in laboratory experiments and behavior in the field. Laboratory research (and, to some extent, also survey research) represents a context in which a rich set of questions can be asked of each individual, and thus a precise estimate of an individual's risk preferences is, in principle, possible. An open question is how useful such estimates are for real-world applications.

In fact, there is research that starts to address this question. Most of this research uses surveys to obtain both a measure of risk aversion and self-reports of field behavior *for the same agents*. For instance, Barsky et al. (1997) use data from the 1992 Health and Retirement Study. This survey included two hypothetical (binary) choices between gambles, framed as whether one would give up a safe job for an alternative job that is risky. Respondents are categorized into four risk groups based on their answers to the two questions. The survey also collects

self-reports on various risky behaviors. Barsky et al. find that the categorical measure of risk preferences is indeed correlated with the likelihood of engaging in risky behaviors such as smoking, drinking, not having insurance, choosing risky employment, and holding risky assets. For similar analyses, see Guiso and Paiella (2006, 2008) using the Bank of Italy's Survey of Household Income and Wealth, and Dohmen et al. (2011) using the German Socio-Economic Panel. Also, in a development context, Liu (2013) and Liu and Huang (2013) study how risk preferences elicited in a (field) laboratory setting correlate, respectively, with the adoption of new crops and the usage of pesticide by Chinese farmers.

An alternative approach is to gather experimental and field data on arguably the same or very similar choices and then investigate the extent to which risk preferences estimated on experimental data correspond to risk preferences estimated on field data. One paper that follows this approach is Barseghyan, O'Donoghue, and Xu (2015). In an online survey, subjects are presented with three deductible choices for property-insurance coverages. The choice menus are constructed to match closely the deductible options and prices associated with specific households in the field data used by BMOT. The key difference is that agents in the laboratory are provided with their latent (objective) claim probabilities, while BMOT assume households' beliefs correspond to observed claim probabilities.

Qualitatively, the findings of Barseghyan, O'Donoghue, and Xu (2015) confirm the patterns documented with field data (BMOT; Barseghyan, Molinari, and Teitelbaum 2016): probability distortions can rationalize the behavior of the vast majority of households, while the curvature of the utility function alone cannot do so. The estimated (average) probability distortion function is increasing in the relevant range and exhibits

significant overweighting. Quantitatively, there are interesting differences between the laboratory and field findings. There is more overweighting in the laboratory than in the field, as well as more heterogeneity.

Going forward, we envision more work emerging in this area, as more field data become available, and researchers are granted ways to design surveys that can reach (a subset of) subjects in the field data (e.g., as in Handel and Kolstad 2015).

7.3 Direct Measurement of Beliefs

When estimating risk preferences from field data, a researcher typically faces a fundamental identification problem: observed choices are often consistent with many combinations of decision makers' risk preferences and their subjective expectations about various outcome probabilities. This problem is typically solved in one of two ways. One approach is to assume individuals hold objective expectations. In the insurance context, this translates into assuming that individuals know the objective claim probabilities or the objective health shocks distribution, and in the betting context this translates into assuming that individuals know horses' objective odds of winning races. The analysis then focuses on identification and estimation of risk preferences. A second approach is to assume that individuals are risk neutral, and the analysis focuses on identification and estimation of subjective beliefs.⁸⁴

A different and promising approach for solving this identification problem is to measure probabilistic expectations directly, as advocated in Manski (2004). For example, one may elicit subjective beliefs on the likelihood of a claim with questions such as: "What

do you think is the percent chance that you will experience an auto collision claim within the next twelve months?" An important aspect of this approach is that expectations are not elicited in qualitative form (e.g., by asking individuals how likely to occur an event is, with options "very," "fairly," "not too," or "not at all" likely), but in quantitative numerical form (e.g., on a scale out of 100).

An extensive literature demonstrates that respondents are willing and able to provide this information in probabilistic format. Manski (2004) and Hurd (2009) survey these findings in connection with data elicited in developed countries, while Delavande (2014) does the same in connection with data elicited in developing countries. The findings include, *inter alia*, that there is substantial heterogeneity in beliefs (at least in the contexts analyzed so far), and that elicited beliefs correlate with individuals' observable characteristics similarly to how actual outcomes do; see, e.g., Dominitz (1998) and Hurd and McGary (2002). While it is not possible to evaluate directly whether the reported expectations are in fact those that respondents truly hold (because there cannot be validation data for this information), numerous studies in different contexts have shown that, when asked about questions that are relevant to their lives, respondents give internally consistent answers. Examples include Dominitz and Manski (1997), who study individuals' income expectations, and Manski and Straub (2000), who analyze individuals' expectations of their job security, in both cases analyzing data from the Survey of Economic Expectations. Manski (2004, sections 5 and 6) summarizes their findings and the findings of many other studies.

For the literature on estimating risk preferences in the field, using data on subjective expectations in place of objective probabilities seems a promising approach for future work. Indeed, in other contexts, probabilistic expectations data have already been

⁸⁴A third approach is to assume more structure in the model of risk preferences. For instance, if one imposes EU, which in particular assumes that whatever decision weights are applied to outcomes must correspond to subjective beliefs, then, with appropriate data, one can identify both the utility function and subjective beliefs.

used to enrich econometric analysis of field data, including retirement behavior (e.g., Hurd, Smith, and Zissimopoulos 2004; van der Klaauw and Wolpin 2008), criminal behavior (e.g., Lochner 2007), contraceptive choices and updating of beliefs on contraceptive effectiveness (e.g., Delavande 2008a, 2008b), and schooling choices (e.g., Giustinelli 2016 and Wiswall and Zafar 2015).⁸⁵

7.4 Assumptions about Mental Accounting

As we discuss in section 5.1.1, in order to estimate risk preferences in a field context, one must make assumptions about how the typically complex field context is translated into concrete choice data that can be used to estimate risk preferences. Such a step is, of course, part of most economic analyses—economics is all about developing simplified models that capture the essence of a field context of interest. However, a second possible justification for such assumptions is that the agents themselves transform—and simplify—a choice situation in their own minds before making a decision. This type of mental operation on the part of agents is often labeled “mental accounting,” and thus we refer to these assumptions as mental accounting assumptions. An important agenda for future research is to pay more careful attention to these assumptions and investigate directly the impact of such assumptions on estimates of risk preferences.

An important dimension on which one must make a mental accounting assumption is how broadly versus narrowly households bracket their decisions. On one extreme, households could bracket all their decisions together into one grand “life” decision—indeed, theoretical economic models are often written in this way. On the other

extreme, households could bracket very narrowly and evaluate each decision in isolation from all others. There are also many possibilities in between.

In fact, virtually all papers that estimate risk preferences implicitly—and occasionally explicitly—assume very narrow bracketing. They estimate risk preferences reflected in one particular choice in isolation from how that choice might interact with the many other choices that households make. In most cases, narrow bracketing is assumed merely to help the researcher. If one has data on only one decision per household, it is hard to assume anything other than narrow bracketing. Even when one has data on multiple decisions per household, it can be computationally burdensome to collect them all together into one grand decision.⁸⁶ There may, however, be some psychological realism to the assumption of narrow bracketing. Indeed, there is a literature that suggests that when people make multiple choices, they frequently do not assess the consequences in an integrated way, but rather tend to make each choice in isolation (e.g., Read, Loewenstein, and Rabin 1999).

In future research, it is worth investigating more carefully how broadly households bracket their decisions. In simple terms, we need to understand whether estimates of risk preferences are biased by incorrect assumptions about bracketing. Even beyond this, if households do in fact bracket multiple decisions together, then they are choosing between more complex lotteries and, as highlighted in section 4, this might permit one to separately identify multiple sources of aversion to risk. Indeed, Barseghyan et al. (2013a) propose exactly this approach.

⁸⁶For instance, if one assumed daily bracketing at a horse track, and if on any day the track holds ten races with eight horses in each, then there are 8^{10} possible lotteries that one could choose (and even this is restricting attention to win bets and ignoring any dynamics associated with basing later bets on earlier results).

⁸⁵Probabilistic expectations data have also been used to enrich experimental analyses, for example, by Nyarko and Schotter (2002) and Dominitz and Hung (2009).

Another aspect of bracketing is that, in many analyses, researchers treat what is really a dynamic choice as a static choice. Perhaps the best example is research on insurance choices. When people choose an insurance deductible, this is really part of a dynamic problem where they are choosing the budget constraints they'll face in the future as a function of intervening events. Framed this way, it becomes clear that a major confound in insurance analyses is a fear of future liquidity constraints. If having to pay a \$1,000 deductible will leave a household with little liquidity to spend on other needs, that household might choose a lower deductible not because of risk aversion with regard to lifetime wealth, but rather because of fear of short-term liquidity issues.⁸⁷ The literature using macroeconomic data performs better on this dimension, as it explicitly incorporates the dynamic life-cycle problems that households face. Nonetheless, even that literature tends to abstract away from household-specific details such as short-run liquidity constraints.

A second dimension on which one must make a mental accounting assumption is what options enter a household's consideration set—that is, what options does a household take to be in its choice set. Because estimates of risk preferences can depend on what options are considered but not chosen, assumptions about the consideration set can alter estimates. To illustrate, consider a stylized example: Suppose we observe an individual in a casino who chooses to bet \$10 on BLACK in roulette (on a typical roulette wheel that has both 0 and 00). This choice yields a risky lottery (+\$10, 18/38; -\$10, 20/38) that has an expected value of -\$0.526.

This person could have also bet \$10 on #1, which instead would yield a riskier lottery (+\$360, 1/38; -\$10, 37/38) that has a larger expected value of -\$0.263. If we estimated, for example, an EU model focusing on the fact that the person chose to bet on BLACK rather than bet on #1, we would conclude that the person is risk averse, and we could infer a lower bound on the magnitude of this risk aversion. However, if instead we estimated an EU model focusing on the fact that the person chose to bet on BLACK rather than to not bet at all, we would instead conclude that the person is risk loving, and we could infer a lower bound on the magnitude of this risk lovingness.

The intuition of this example extends almost immediately to research that estimates risk preferences using data from horse races. In such analyses, one must make an assumption about whether the option not to bet is included in the consideration set. Because the expected return on most horses is negative, including the option not to bet in the consideration set will yield estimates of risk preferences that are more risk loving.⁸⁸ But this issue also applies to property insurance. When estimating preferences from deductible choices, one might wonder whether households consider the possibility of not insuring at all, and also whether they turn down any insurance options from other firms that are not in the data set. In future research, it is worth investigating more carefully the determinants and importance of consideration sets.

A third dimension on which one must make a mental accounting assumption is what are the possible outcomes that households consider. To illustrate, consider property insurance. When researchers estimate risk

⁸⁷ Even relatively wealthy households can face short-run liquidity concerns if they hold a large share of their wealth in illiquid assets. See Kaplan, Violante, and Weidner (2014) for evidence on the prevalence of such “wealthy hand-to-mouth” households.

⁸⁸ This issue is mostly solved if one estimates utility nonparametrically, as in Snowberg and Wolfers (2010), where the estimated utility for not betting has no impact on the estimated utility for winning or losing money.

preferences using data on property insurance, they typically assume that households only consider the possibility of incurring no loss or a single loss during the policy period. However, in principle, one might incur two, three, or even more losses during a policy period. If so, then the set of lotteries from which households are choosing are different. The assumption of zero or one loss is often made for simplicity, but again it could reflect a psychological realism, as people seem to have a hard time imagining and accounting for all the possibilities that could occur in life. This is especially true in more complex domains, such as health insurance, where it seems quite likely that households approach decisions with a simplified conceptualization of all the possible outcomes that might occur.

Moving forward, we think it important that the literature considers more carefully and more directly these and other mental accounting assumptions when estimating risk preferences. Such assumptions can matter under EU, and they become even more important under RDEU and other more complex models of risk preferences.

7.5 Concluding Thoughts

The literature on identification and estimation of risk preferences using field data has improved a lot over the years, especially in the past decade. As more and more forms of field data become available, we expect the literature will continue to blossom, and we hope our analysis in this review will help frame the development of the literature.

In particular, we have highlighted a number of important dimensions on which researchers must make some key choices. In looking for field contexts, one must think through the trade-offs associated with individual versus aggregate data, as well as those associated with studying simple versus complex choices. In developing a model of how individuals make choices,

researchers must make three broad categories of assumptions: (i) assumptions about how the modeler—and perhaps the agents themselves—transform and simplify the real-world choice situation into a well-defined choice; (ii) assumptions about agents' subjective beliefs, the extent to which they correspond to objective beliefs, and the extent of heterogeneity; and (iii) assumptions about the possible sources of aversion to risk. Moreover, when one uses aggregate data, there is a fourth category of assumptions: (iv) assumptions about the nature of market forces. Our analysis has highlighted the importance of all of these dimensions, and indeed, our discussion of directions for future research in this section suggests yet new ways to approach many of these dimensions. There clearly is much to be done, and we look forward to watching it all evolve.

REFERENCES

- Abel, Andrew B. 1990. "Asset Prices under Habit Formation and Catching Up with the Joneses." *American Economic Review* 80 (2): 38–42.
- Abel, Andrew B. 1999. "Risk Premia and Term Premia in General Equilibrium." *Journal of Monetary Economics* 43 (1): 3–33.
- Ali, Mukhtar M. 1977. "Probability and Utility Estimates for Racetrack Bettors." *Journal of Political Economy* 85 (4): 803–15.
- Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström. 2008. "Risk Aversion in Game Shows." In *Risk Aversion in Experiments*, edited by James C. Cox and Glenn W. Harrison, 359–404. Bingley: Emerald, JAI Press.
- Apesteguia, Jose, and Miguel A. Ballester. 2018. "Monotone Stochastic Choice Models: The Case of Risk and Time Preferences." *Journal of Political Economy* 126 (1): 74–106.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance." *Quarterly Journal of Economics* 130 (4): 1623–67.
- Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. 2013a. "Distinguishing Probability Weighting from Risk Misperceptions in Field Data." *American Economic Review* 103 (3): 580–85.
- Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. 2013b. "The Nature of Risk Preferences: Evidence from Insurance Choices." *American Economic Review* 103 (6): 2499–529.

- Barseghyan, Levon, Francesca Molinari, and Joshua C. Teitelbaum. 2016. "Inference under Stability of Risk Preferences." *Quantitative Economics* 7 (2): 367–409.
- Barseghyan, Levon, Ted O'Donoghue, and Lin Xu. 2015. "How Different Are Insurance Purchases in Experiments and the Real World?" Unpublished.
- Barseghyan, Levon, Jeffrey Prince, and Joshua C. Teitelbaum. 2011. "Are Risk Preferences Stable across Contexts? Evidence from Insurance Data." *American Economic Review* 101 (2): 591–631.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics* 112 (2): 537–79.
- Beetsma, Roel M. W. J., and Peter C. Schotman. 2001. "Measuring Risk Attitudes in a Natural Experiment: Data from the Television Game Show Lingo." *Economic Journal* 111 (474): 821–48.
- Bell, David E. 1985. "Disappointment in Decision Making under Uncertainty." *Operations Research* 33 (1): 1–27.
- Bernheim, B. Douglas. 2009. "Behavioral Welfare Economics." *Journal of the European Economic Association* 7 (2–3): 267–319.
- Berry, Steven, Amit Gandhi, and Philip Haile. 2013. "Connected Substitutes and Invertibility of Demand." *Econometrica* 81 (5): 2087–111.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–90.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Options." *Quarterly Journal of Economics*, 132 (3): 1319–72.
- Bombardini, Matilde, and Francesco Trebbi. 2012. "Risk Aversion and Expected Utility Theory: An Experiment with Large and Small Stakes." *Journal of the European Economic Association* 10 (6): 1348–99.
- Botti, Fabrizio, Anna Conte, Daniela T. Di Cagno, and Carlo D'Ippoliti. 2008. "Risk Attitude in Real Decision Problems." *B.E. Journal of Economic Analysis and Policy* 8 (1).
- Camerer, Colin F. 1995. "Individual Decision Making." In *Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth, 587–703. Princeton and Oxford: Princeton University Press.
- Camerer, Colin F. 2000. "Prospect Theory in the Wild: Evidence from the Field." In *Choices, Values, and Frames*, edited by Daniel Kahneman and Amos Tversky, 288–300. Cambridge and New York: Cambridge University Press.
- Campbell, John Y. 2003. "Consumption-Based Asset Pricing." In *Handbook of the Economics of Finance: Volume 1B*, edited by George M. Constantinides, Milton Harris, and Rene M. Stulz, 803–87. Amsterdam: Elsevier.
- Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–51.
- Cardon, James H., and Igal Hendel. 2001. "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditures Survey." *RAND Journal of Economics* 32 (3): 408–27.
- Chetty, Raj. 2006. "A New Method of Estimating Risk Aversion." *American Economic Review* 96 (5): 1821–34.
- Chetty, Raj. 2015. "Behavioral Economics and Public Policy: A Pragmatic Perspective." *American Economic Review* 105 (5): 1–33.
- Chiappori, Pierre-André, Bernard Salanié, François Salanié, and Amit Gandhi. 2012. "From Aggregate Betting Data to Individual Risk Preferences." Columbia University Department of Economics Discussion Paper No. 1213-08.
- Cicchetti, Charles J., and Jeffrey A. Dubin. 1994. "A Microeconomic Analysis of Risk Aversion and the Decision to Self-Insure." *Journal of Political Economy* 102 (1): 169–86.
- Cohen, Alma, and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review* 97 (3): 745–88.
- Constantinides, George M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98 (3): 519–43.
- Conte, Anna, Peter G. Moffatt, Fabrizio Botti, Daniela T. Di Cagno, and Carlo D'Ippoliti. 2012. "A Test of the Rational Expectations Hypothesis Using Data from a Natural Experiment." *Applied Economics* 44 (35): 4661–78.
- de Roos, Nicolas, and Yianis Sarafidis. 2010. "Decision Making Under Risk in *Deal or No Deal*." *Journal of Applied Econometrics* 25 (6): 987–1027.
- Deck, Cary, Jungmin Lee, and Javier Reyes. 2008. "Risk Attitudes in Large Stake Gambles: Evidence from a Game Show." *Applied Economics* 40 (1): 41–52.
- Delavande, Adeline. 2008a. "Measuring Revisions to Subjective Expectations." *Journal of Risk and Uncertainty* 36 (1): 43–82.
- Delavande, Adeline. 2008b. "Pill, Patch, or Shot? Subjective Expectations and Birth Control Choice." *International Economic Review* 49 (3): 999–1042.
- Delavande, Adeline. 2014. "Probabilistic Expectations in Developing Countries." *Annual Review of Economics* 6: 1–20.
- Diecidue, Enrico, and Peter P. Wakker. 2001. "On the Intuition of Rank-Dependent Utility." *Journal of Risk and Uncertainty* 23 (3): 281–98.
- Dohmen, Thomas, David Huffman, Jürgen Schupp, Armin Falk, Uwe Sunde, and Gert G. Wagner. 2011. "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences." *Journal of the European Economic Association* 9 (3): 522–50.
- Dominitz, Jeff. 1998. "Earnings Expectations, Revisions, and Realizations." *Review of Economics and Statistics* 80 (3): 374–88.
- Dominitz, Jeff, and Angela A. Hung. 2009. "Empirical Models of Discrete Choice and Belief Updating in Observational Learning Experiments." *Journal of Economic Behavior and Organization* 69 (2): 94–109.

- Dominitz, Jeff, and Charles F. Manski. 1997. "Using Expectations Data to Study Subjective Income Expectations." *Journal of the American Statistical Association* 92 (439): 855–67.
- Edwards, Ward. 1955. "The Prediction of Decisions among Bets." *Journal of Experimental Psychology* 50: 201–14.
- Edwards, Ward. 1962. "Subjective Probabilities Inferred from Decisions." *Psychological Review* 69 (2): 109–35.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877–921.
- Einav, Liran, Amy Finkelstein, Iuliana Pascu, and Mark R. Cullen. 2012. "How General Are Risk Preferences? Choices under Uncertainty in Different Domains." *American Economic Review* 102 (6): 2606–38.
- Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. "Selection on Moral Hazard in Health Insurance." *American Economic Review* 103 (1): 178–219.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2010. "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market." *Econometrica* 78 (3): 1031–92.
- Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4): 937–69.
- Epstein, Larry G., and Stanley E. Zin. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy* 99 (2): 263–86.
- Ericson, Keith Marzilli, Philipp Kircher, Johannes Spinnewijn, and Amanda Starc. 2018. "Inferring Risk Perceptions and Preferences Using Choice from Insurance Menus: Theory and Evidence." Unpublished.
- Fehr-Duda, Helga, and Thomas Epper. 2012. "Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences." *Annual Review of Economics* 4: 567–93.
- Fullenkamp, Connel, Rafael Tenorio, and Robert Battalio. 2003. "Assessing Individual Risk Attitudes Using Field Data from Lottery Games." *Review of Economics and Statistics* 85 (1): 218–26.
- Gandhi, Amit, and Ricardo Serrano-Padial. 2015. "Does Belief Heterogeneity Explain Asset Prices: The Case of the Longshot Bias." *Review of Economic Studies* 82 (1): 156–86.
- Gertner, Robert. 1993. "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'." *Quarterly Journal of Economics* 108 (2): 507–21.
- Giustinelli, Pamela. 2016. "Group Decision Making with Uncertain Outcomes: Unpacking Child-Parent Choice of the High School Track." *International Economic Review* 57 (2): 573–602.
- Gonzalez, Richard, and George Wu. 1999. "On the Shape of the Probability Weighting Function." *Cognitive Psychology* 38 (1): 129–66.
- Griffith, Richard M. 1949. "Odds Adjustments By American Horse-Race Bettors." *American Journal of Psychology* 62 (2): 290–94.
- Griffith, Richard M. 1961. "A Footnote on Horse Race Betting." *Transactions, Kentucky Academy of Science* 22: 78–81.
- Guiso, Luigi, and Monica Paiella. 2006. "Competitive Failures in Insurance Markets: Theory and Policy Implications." In *Competitive Failures in Insurance Markets: Theory and Policy Implications*, edited by Pierre-Andre Chiappori and Christian Gollier, 213–50. Cambridge and London: MIT Press.
- Guiso, Luigi, and Monica Paiella. 2008. "Risk Aversion, Wealth, and Background Risk." *Journal of the European Economic Association* 6 (6): 1109–50.
- Gul, Faruk. 1991. "A Theory of Disappointment Aversion." *Econometrica* 59 (3): 667–86.
- Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.
- Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–2500.
- Hansen, Lars Peter, and Kenneth J. Singleton. 1982. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica* 50 (5): 1269–86.
- Hansen, Lars Peter, and Kenneth J. Singleton. 1983. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy* 91 (2): 249–65.
- Hartley, Roger, Gauthier Lanot, and Ian Walker. 2014. "Who Really Wants to Be a Millionaire? Estimates of Risk Aversion from Gameshow Data." *Journal of Applied Econometrics* 29 (6): 861–79.
- Harville, David A. 1973. "Assigning Probabilities to the Outcomes of Multi-entry Competitions." *Journal of the American Statistical Association* 68 (342): 312–16.
- Hurd, Michael D. 2009. "Subjective Probabilities in Household Surveys." *Annual Review of Economics* 1: 543–62.
- Hurd, Michael D., and Kathleen McGarry. 2002. "The Predictive Validity of Subjective Probabilities of Survival." *Economic Journal* 112 (482): 966–85.
- Hurd, Michael D., James P. Smith, and Julie M. Zissimopoulos. 2004. "The Effects of Subjective Survival on Retirement and Social Security Claiming." *Journal of Applied Econometrics* 19 (6): 761–75.
- Jullien, Bruno, and Bernard Salanié. 2000. "Estimating Preferences under Risk: The Case of Racetrack Bettors." *Journal of Political Economy* 108 (3): 503–30.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–92.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*: 77–138.

- Karmarkar, Uday S. 1978. "Subjectively Weighted Utility: A Descriptive Extension of the Expected Utility Model." *Organizational Behavior and Human Performance* 21 (1): 61–72.
- Kocherlakota, Narayana. 1996. "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature* 34 (1): 42–71.
- Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65.
- Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-Dependent Risk Attitudes." *American Economic Review* 97 (4): 1047–73.
- Kőszegi, Botond, and Matthew Rabin. 2008. "Choices, Situations, and Happiness." *Journal of Public Economics* 92 (8–9): 1821–32.
- Lattimore, Pamela K., Joanna R. Baker, and Ann D. Witte. 1992. "The Influence of Probability on Risky Choice: A Parametric Examination." *Journal of Economic Behavior and Organization* 17 (3): 377–400.
- Lewbel, Arthur. 2017. "The Identification Zoo—Meanings of Identification in Econometrics." Unpublished.
- Liu, Elaine M. 2013. "Time to Change What to Sow: Risk Preferences and Technology Adoption Decisions of Cotton Farmers in China." *Review of Economics and Statistics* 95 (4): 1386–403.
- Liu, Elaine M., and JiKun Huang. 2013. "Risk Preferences and Pesticide Use By Cotton Farmers in China." *Journal of Development Economics* 103: 202–15.
- Lochner, Lance. 2007. "Individual Perceptions of the Criminal Justice System." *American Economic Review* 97 (1): 444–60.
- Loomes, Graham, and Robert Sugden. 1986. "Disappointment and Dynamic Consistency in Choice under Uncertainty." *Review of Economic Studies* 53 (2): 271–82.
- Manski, Charles F. 1977. "The Structure of Random Utility Models." *Theory and Decision* 8 (3): 229–54.
- Manski, Charles F. 2004. "Measuring Expectations." *Econometrica* 72 (5): 1329–76.
- Manski, Charles F., and John D. Straub. 2000. "Worker Perceptions of Job Insecurity in the Mid-1990s: Evidence from the Survey of Economic Expectations." *Journal of Human Resources* 35 (3): 447–79.
- Markowitz, Harry. 1952. "The Utility of Wealth." *Journal of Political Economy* 60 (2): 151–58.
- Masatlioglu, Yusufcan, and Collin Raymond. 2016. "A Behavioral Analysis of Stochastic Reference Dependence." *American Economic Review* 106 (9): 2760–82.
- McFadden, Daniel. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Paul Zarembka, 105–42. New York: Academic Press.
- McGlothlin, William H. 1956. "Stability of Choices among Uncertain Alternatives." *American Journal of Psychology* 69 (4): 604–15.
- Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15: 145–61.
- Metrick, Andrew. 1995. "A Natural Experiment in 'Jeopardy!'" *American Economic Review* 85 (1): 240–53.
- Mulino, Daniel, Richard Scheelings, Robert Brooks, and Robert Faff. 2009. "Does Risk Aversion Vary with Decision-Frame? An Empirical Test Using Recent Game Show Data." *Review of Behavioral Finance* 1 (1–2): 44–61.
- Nyarko, Yaw, and Andrew Schotter. 2002. "An Experimental Study of Belief Learning Using Elicited Beliefs." *Econometrica* 70 (3): 971–1005.
- Post, Thierry, Martijn van den Assem, Guido Baltussen, and Richard H. Thaler. 2008. "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show." *American Economic Review* 98 (1): 38–71.
- Pratt, John W. 1964. "Risk Aversion in the Small and in the Large." *Econometrica* 32 (1): 122–36.
- Prelec, Drazen. 1998. "The Probability Weighting Function." *Econometrica* 66 (3): 497–527.
- Preston, Malcolm G., and Philip Baratta. 1948. "An Experimental Study of the Auction-Value of an Uncertain Outcome." *American Journal of Psychology* 61 (2): 183–93.
- Quiggin, John. 1982. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization* 3 (4): 323–43.
- Rabin, Matthew. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68 (5): 1281–92.
- Rabin, Matthew, and Richard H. Thaler. 2001. "Anomalies: Risk Aversion." *Journal of Economic Perspectives* 15 (1): 219–32.
- Read, Daniel, George Loewenstein, and Matthew Rabin. 1999. "Choice Bracketing." *Journal of Risk and Uncertainty* 19 (1–3): 171–97.
- Rees-Jones, Alex. 2018. "Quantifying Loss-Averse Tax Manipulation." *Review of Economic Studies* 85 (2): 1251–78.
- Siegel, Jeremy J., and Richard H. Thaler. 1997. "Anomalies: The Equity Premium Puzzle." *Journal of Economic Perspectives* 11 (1): 191–200.
- Snowberg, Erik, and Justin Wolfers. 2010. "Explaining the Favorite–Long Shot Bias: Is it Risk-Love or Misperceptions?" *Journal of Political Economy* 118 (4): 723–46.
- Starc, Amanda. 2014. "Insurer Pricing and Consumer Welfare: Evidence from Medigap." *RAND Journal of Economics* 45 (1): 198–220.
- Starmer, Chris. 2000. "Developments in Non-expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk." *Journal of Economic Literature* 38 (2): 332–82.
- Sydnor, Justin. 2010. "(Over)insuring Modest Risks." *American Economic Journal: Applied Economics* 2 (4): 177–99.
- Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.
- van der Klaauw, Wilbert, and Kenneth I. Wolpin. 2008. "Social Security and the Retirement and Savings

- Behavior of Low-Income Households." *Journal of Econometrics* 145 (1–2): 21–42.
- Weil, Philippe. 1989. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *Journal of Monetary Economics* 24 (3): 401–21.
- Weil, Philippe. 1990. "Nonexpected Utility in Macroeconomics." *Quarterly Journal of Economics* 105 (1): 29–42.
- Weitzman, Martin. 1965. "Utility Analysis and Group Behavior: An Empirical Study." *Journal of Political Economy* 73 (1): 18–26.
- Wiswall, Matthew, and Basit Zafar. 2015. "Determinants of College Major Choice: Identification Using an Information Experiment." *Review of Economic Studies* 82 (2): 791–824.
- Wolf, Charles, and Larry Pohlman. 1983. "The Recovery of Risk Preferences from Actual Choices." *Econometrica* 51 (3): 843–50.
- Yaari, Menahem E. 1965. "Convexity in the Theory of Choice under Risk." *Quarterly Journal of Economics* 79 (2): 278–90.