

Household Portfolio Underdiversification and Probability Weighting: Evidence from the Field

Stephen G. Dimmock, Roy Kouwenberg, Olivia S. Mitchell, and Kim Peijnenburg

August 8, 2018

PRC WP2018-9

Pension Research Council Working Paper Pension Research Council

The Wharton School, University of Pennsylvania
3620 Locust Walk, 3000 SH-DH
Philadelphia, PA 19104-6302
Tel.: 215.573.3414 Fax: 215.573.3418
Email: prc@wharton.upenn.edu
<http://www.pensionresearchcouncil.org>

This paper is part of the NBER's Research Program on the Economics of Aging and the Working Group on Household Portfolios. For assistance with the survey, the authors gratefully acknowledge Tim Colvin, David Grant, and the American Life Panel team at the RAND Corporation. For comments, we thank Nishad Kapadia, Alexander Klos, Sonya Lim, Gianpaolo Parise, Valery Polkovnichenko, Peter Wakker, and Joe Zhang; seminar participants at EDHEC, HEC Paris, IESE Business School, London School of Economics, University of Cambridge, and University of Warwick; and participants at the Financial Intermediation Research Society, SFS Cavalcade, Singapore Symposium, TIAA/PRC Symposium, and Western Finance Association. Yunju Cha and Yong Yu provided outstanding research assistance. This project received funding from the TIAA Institute and Wharton School's Pension Research Council/Boettner Center. The content is solely the responsibility of the authors and does not represent official views of the TIAA Institute or the Wharton School's Pension Research Council/Boettner Center. Online Appendices are at: https://pensionresearchcouncil.wharton.upenn.edu/wp-content/uploads/2018/08/Online_Appendix_2018-08-08.pdf ©2018 Dimmock, Kouwenberg, Mitchell, and Peijnenburg.

Household Portfolio Underdiversification and Probability Weighting: Evidence from the Field

Stephen G. Dimmock, Roy Kouwenberg, Olivia S. Mitchell, and Kim Peijnenburg

Abstract

We explore the relation between probability weighting and household portfolio underdiversification in a representative household survey, using custom-designed incentivized lotteries. On average, people display Inverse-S shaped probability weighting, overweighting the small probabilities of tail events. As theory predicts, our Inverse-S measure is positively associated with portfolio underdiversification, which results in significant Sharpe ratio losses. We match respondents' individual stock holdings to CRSP data and find that people with higher Inverse-S tend to pick stocks with positive skewness and hold positively-skewed equity portfolios. We show that these choices reflect preferences rather than probability unsophistication or limited financial knowledge.

JEL Codes: G11, D81, D14, C83

Keywords: household finance, portfolio underdiversification, probability weighting, rank dependent utility, cumulative prospect theory, household portfolio puzzles, stock market participation

Stephen G. Dimmock
Nanyang Technological University
dimmock@ntu.edu.sg

Roy Kouwenberg
Mahidol University and Erasmus University
Rotterdam
roy.kou@mahidol.ac.th

Olivia S. Mitchell
The Wharton School of the University of
Pennsylvania and NBER
mitchelo@wharton.upenn.edu

Kim Peijnenburg
HEC Paris and CEPR
peijnenburg@hec.fr

Household Portfolio Underdiversification and Probability Weighting: Evidence from the Field

People frequently violate the tenets of expected utility theory for low probability events: for example, they simultaneously buy insurance and lottery tickets, overinsure against small losses, and hold undiversified positions in individual company stocks with high positive skewness hoping to pick the “next Apple.”¹ Such seemingly anomalous behaviors are consistent with *probability weighting*: the idea that people use transformed rather than objective probabilities when making decisions. As formalized in prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) and rank-dependent utility theory (Quiggin, 1982; Yaari, 1987), people tend to overweight low probability tail events and underweight events from the middle of the probability distribution.

Several theoretical papers show that probability weighting predicts anomalies in observed household portfolio decisions, such as underdiversification and the popularity of lottery-type stocks (e.g., Barberis and Huang, 2008). For individuals who overweight small probability tail events, the negative skewness of the aggregate stock market makes a well-diversified portfolio less attractive; but the positive skewness of an underdiversified portfolio containing a few individual stocks becomes more attractive.²

The empirical literature is less developed, offering mostly indirect evidence using calibrated portfolio choice models (e.g., Polkovnichenko, 2005). Obtaining direct empirical evidence on the role of probability weighting is challenging, because individual preferences such as probability weighting are not readily observable. The present paper provides direct evidence

¹ For further discussion, see the review articles of Fehr-Duda and Epper (2012) and Barberis (2013a).

² See Shefrin and Statman (2000), Polkovnichenko (2005), Barberis and Huang (2008), Jin and Zhou (2008), Chapman and Polkovnichenko (2011), De Giorgi and Legg (2012), and He, Kouwenberg, and Zhou (2018). Albuquerque (2012) presents a theoretical foundation for the positive skewness of individual stocks.

that probability weighting can explain actual household portfolio decisions, most notably portfolio underdiversification, skewness seeking, and investments in lottery-type stocks.

To elicit individuals' probability weighting preferences, we designed a purpose-built internet survey module and fielded it in a nationally-representative sample of several thousand respondents in the American Life Panel (ALP). Our module elicits certainty equivalents for a series of binary lotteries adapted from Wakker and Deneffe (1996) and Abdellaoui (2000). The probabilities of winning the lotteries vary from small to large, allowing us to obtain a non-parametric measure of individual respondents' probability weighting behavior which we term *Inverse-S*. In addition to a fixed participation fee, all respondents had the opportunity to receive real monetary incentives based on their choices (we paid a total of \$16,020 to 2,072 of the 2,703 eligible respondents). The survey module also measures subjects' portfolio allocations and collects the names of their five largest individual stockholdings.

Our general population estimates of probability weighting are consistent with those found in laboratory studies (Abdellaoui, 2000; Booi, van Praag, and van de Kuilen, 2010; Bruhin, Fehr-Duda, and Epper, 2010). Specifically, we show that most people have *Inverse-S* shaped probability weighting functions implying overweighting of tail events, though there is substantial cross-subject heterogeneity. On average, when the probability of winning a lottery is only 5%, our subjects' certainty equivalent is *greater* than the expected value of the lottery, which is consistent with overweighting the small probability of winning. By contrast, when the probability of winning a lottery is higher (e.g., 50%), our subjects' certainty equivalent is *less* than the expected value of the lottery.

Using our subject-specific variable, *Inverse-S*, we test the theoretical predictions regarding probability weighting and portfolio choice. Specifically, we explore the portfolios of equity

holders and measure the fraction of total equity allocated to individual stocks, which Calvet, Campbell, and Sodini (2007, 2009) show is a good proxy for portfolio underdiversification.³ We find that a one standard deviation increase in *Inverse-S* implies a 12.7 percentage point increase in the fraction of the portfolio allocated to individual stocks (28.2% relative to the baseline allocation of 45.0 percentage points). Using subjects' individual stock holdings, we construct an alternative measure of underdiversification; the relative Sharpe ratio loss from investing in individual stocks (as compared to investing in the market portfolio; see Calvet, Campbell, and Sodini, 2007). We find that high *Inverse-S* is associated with large Sharpe ratio losses due to idiosyncratic risk. In particular, our results imply that a one-standard deviation higher *Inverse-S* implies a cost to the average (median) stockholder of \$2,504 (\$351) per year, as for the same amount of risk the person could have had a higher expected return.

Our results are robust to controlling for variables commonly used in the literature such as age, income, financial wealth, education, marital status, number of household members, and employment, as well as additional controls for subjects' risk aversion, financial literacy, trust, optimism, and numeracy. Additionally, the module recorded the time each subject spent on the elicitation questions and included check questions to assess whether subjects' choices were internally consistent. Our results are also robust to excluding subjects who answered the elicitation questions unusually quickly or who made multiple errors on the check questions. Moreover, we find similar results when using alternative probability weighting measures based on specific functional forms (e.g., the functions proposed by Prelec, 1998, and Bordalo, Gennaioli, and Shleifer, 2012).

³ Consistent with Calvet, Campbell, and Sodini (2007, 2009), we find that, conditional on owning individual stocks, half of the respondents held shares in only one or two individual companies, which is consistent with the fraction of the equity portfolio allocated to individual stocks being a reasonable proxy for underdiversification.

In addition, probability weighting can help explain the *type* of individual stocks people choose. To this end, we asked respondents who own individual stocks to provide the names (or tickers) of their five largest holdings. We then match these names to the CRSP daily stock return database and construct various measures of the stocks' characteristics. Consistent with the predictions of theory, we find that respondents with high *Inverse-S* tend to hold individual stocks with high positive (expected) skewness.

Next, we evaluate whether probability weighting is a component of preferences or a symptom of probability unsophistication. For example, poor quantitative reasoning ability could cause both probability weighting and observed portfolio choices. Based on our summary statistics this seems unlikely, as probability weighting is weakly *positively* correlated with education, numerical reasoning ability, and financial literacy. Robustness tests show our results are similar if we restrict the sample to subjects who score higher on proxies for probability sophistication: having a college degree, and making no errors on numerical reasoning questions. The results are also similar if we restrict the sample to subjects who understand that owning a single company's stock is generally riskier than owning a stock mutual fund, or making no errors on all financial literacy questions. Overall, we conclude that probability weighting reflects preferences, and not probability unsophistication or limited financial knowledge.

We then broaden the sample of survey respondents to consider non-participation in equity markets as well as the type of equity held by those who do participate. The theoretical literature shows that probability weighting can result in non-participation due to first-order risk aversion,⁴ and if individuals with high probability weighting do participate they will hold underdiversified individual stock portfolios. To test these predictions, we use a multinomial logit model with four

⁴ See Epstein and Zin (1990), Segal and Spivak (1990), and Chapman and Polkovnichenko (2011).

categories: non-participation, mutual funds only, individual stocks only, and both mutual funds and individual stocks. We find that *Inverse-S* is positively associated with non-participation *and* ownership of individual stocks, and thus it is negatively associated with owning only mutual funds. This result provides evidence that probability weighting is not simply a proxy for risk aversion, as the subjects make *either* the least risky choice (non-participation) or the riskiest choice (an undiversified portfolio): this is inconsistent with predictions following from risk aversion.

Our paper contributes to the household portfolio choice literature by testing theoretical models of probability weighting and household portfolios.⁵ Specifically, it is the first to show a direct relation between elicited probability weighting preferences and actual household portfolio decisions. Relatedly, Polkovnichenko (2005) uses stock return data to obtain numerical results in a calibrated model that links probability weighting and underdiversification. Rieger (2012) and Erner, Klos, and Langer (2013) associate elicited probability weighting metrics to hypothetical financial decisions about structured products in laboratory experiments using university students. In contrast, we relate preferences elicited in the field to people's actual financial decisions. Consistent with the predictions of theory, we show that probability weighting can explain portfolio underdiversification, skewness seeking, and investments in lottery-type stocks, and we further show that probability weighting is related to non-participation in the equity markets.

Our paper also contributes to the empirical literature on household portfolio underdiversification (e.g., Blume and Friend, 1975; Kelly, 1995; Calvet, Campbell, and Sodini, 2007; Mitton and Vorkink, 2007; Goetzmann and Kumar, 2008). For instance, Kumar (2009) finds that households hold underdiversified portfolios and this behavior is related to the demand for

⁵ For example, see Shefrin and Statman (2000), Polkovnichenko (2005), Barberis and Huang (2008), Jin and Zhou (2008), Chapman and Polkovnichenko (2011), De Giorgi and Legg (2012), and He, Kouwenberg, and Zhou (2018), among others.

stocks with lottery-like features. Our paper complements these studies, as it explores the underlying preferences driving this demand.

Finally, our work relates to a branch of the asset pricing literature which posits that probability weighting can explain the historically low returns of many securities with positive skewness. Several authors find that stocks with positive expected skewness have unusually low returns (e.g., Boyer, Mitton, and Vorkink, 2010; Bali, Cakici, and Whitelaw, 2011; Conrad, Dittmar, and Ghysels, 2013; Conrad, Kapadia, and Xing, 2014). Boyer and Vorkink (2014) and Li, Subrahmanyam, and Yang (2018) find similar results for equity options. Though our paper does not directly address asset pricing implications, our findings do support the preference-based explanation offered in the cited studies. That is, we find a direct link between investors' probability weighting preferences and skewness-seeking behavior.

Our paper is the first non-laboratory analysis to provide direct evidence relating probability weighting to households' portfolio choices, specifically, to underdiversification and investment in individual stocks with high positive skewness. These results are consistent with theoretical models predicting that probability weighting can explain puzzling features of household portfolio choice.

1. Eliciting Individuals' Probability Weighting and Utility Curvature

1.1. Rank-Dependent Utility and Probability Weighting

A large body of experimental studies finds that individuals frequently make decisions that contradict the predictions of expected utility (e.g., Camerer, 1995; Starmer, 2000). In the expected utility model, the utility $U(c_i)$ of each outcome c_i is weighted linearly by its probability p_i :

$$E(U) = \sum_{i=1}^N p_i \cdot U(c_i) . \tag{1}$$

However, Allais (1953) demonstrates that linearity in probabilities is often violated in simple choice problems. For example, consider the choice between a 100% certainty of receiving \$1 million versus a 98% chance of winning \$5 million. Most people prefer to receive \$1 million with certainty. Next, consider a modification of this choice in which both probabilities are divided by 100: that is, consider now the choice between a 1% chance of winning \$1 million versus a 0.98% chance of winning \$5 million. Now, most people prefer a 0.98% chance of winning \$5 million. Such a combination of choices is inconsistent with expected utility: the first preference implies $U(1,000,000) > 0.98 \times U(5,000,000)$, while the second preference implies $0.01 \times U(1,000,000) < 0.0098 \times U(5,000,000)$.

This phenomenon, known as the Allais paradox, demonstrates that risk preferences can depend non-linearly on outcome probabilities. Many studies have replicated this finding, including in experiments with large real monetary rewards (e.g., Starmer, 2000). Generally, in both experiments and real world situations, people are risk-seeking when the probability of winning is small but risk averse when the probability of winning is large. Furthermore, many people are risk-seeking for small probabilities of winning but simultaneously risk averse for small probabilities of losing. For example, the same person may buy both lottery tickets and insurance (for a review see Fehr-Duda and Epper, 2012).

A large theoretical and empirical literature shows that Allais' findings can be explained by non-expected utility models in which decision-makers transform probabilities with a non-linear weighting function (Starmer, 2000; Fehr-Duda and Epper, 2012). The two most commonly used models are rank-dependent utility (RDU) developed by Quiggin (1982) and cumulative prospect theory (CPT) developed by Tversky and Kahneman (1992). In these models, individuals order the possible outcomes from worst to best ($c_1 < c_2 < \dots < c_N$) and then assign each outcome a

decision weight π_i that depends on the cumulative probability of the outcome. For example, the utility functional can be given by:

$$V = \sum_{i=1}^N \pi_i \cdot U(c_i) \quad , \quad (2)$$

$$\pi_i = w(P_i) - w(P_{i-1}) = w(p_1 + p_2 + \dots + p_i) - w(p_1 + p_2 + \dots + p_{i-1}) \quad , \quad (3)$$

where π_i is determined by an increasing and differentiable weighting function $w(P_i)$, such that $w(0) = 0$ and $w(1) = 1$, and $P_i = p_1 + p_2 + \dots + p_i$ is the cumulative probability of outcome i .

Figure 1 displays the inverse-S shaped pattern of $w(P_i)$ typically found in experimental studies, in which low probability tail outcomes are substantially overweighted relative to objective probabilities ($\pi_i > p_i$). The weighting function is steep on both the left and the right sides of the figure, which implies that low probability tail outcomes are substantially overweighted for both extreme good outcomes and extreme bad outcomes. This, in turn, can generate risk-seeking towards good outcomes with low probabilities and extreme risk aversion towards bad outcomes with low probabilities.

[Insert Figure 1 here]

Probability weighting is similar in CPT and RDU – the differences between them come from their treatment of utility curvature and not probability weighting – except that in CPT the probabilities for loss outcomes ($c_N < \dots < c_{k+2} < c_{k+1} < 0$) and gain outcomes ($0 < c_k < \dots < c_2 < c_1$) are transformed by two distinct weighting functions, $w^-(P_i)$ and $w^+(P_i)$. Empirically, the weighting functions for losses and gains tend to have the same inverse-S shaped pattern as in Figure 1 (Tversky and Kahneman, 1992). Therefore, low probability outcomes in both tails are overweighted, similar to rank dependent utility.

More recently, Bordalo, Gennaioli, and Shleifer (2012) develop a model in which probability weighting is determined by the salience of the payoffs, with the contrast between payoffs determining their salience. In this model, people overweight the probability of salient gains (losses), resulting in risk-seeking (averse) behavior. Although in some contexts this model generates different predictions than RDU or CPT, for financial choices the predictions are similar. Bordalo, Gennaioli, and Shleifer (2013) show that, relative to expected utility theory, salience theory implies a strong preference for positively skewed securities and reduced demand for a diversified portfolio such as the market index. Accordingly, in this paper we do not seek to distinguish between RDU, CPT, and salience theory.

1.2. Hypotheses

The extant theoretical literature shows that probability weighting affects portfolio choice through two channels: greater sensitivity to skewness and first-order risk aversion. Probability weighting increases sensitivity to skewness, because the investor overweights low probability tail outcomes. Given the negative skewness of the aggregate stock market (Albuquerque, 2012), probability weighting makes owning a well-diversified equity portfolio less attractive (e.g., Polkovnichenko, 2005; Chapman and Polkovnichenko, 2011; De Giorgi and Legg, 2012). In contrast, given the positive skewness of individual stocks, probability weighting makes owning an underdiversified portfolio of individual stocks more attractive (e.g., Shefrin and Statman, 2000; Polkovnichenko, 2005; Barberis and Huang, 2008; Jin and Zhou, 2008). Thus, the prior literature suggests that higher probability weighting will be associated with underdiversification.

We illustrate this prediction using a simple calibrated portfolio choice model with probability weighting preferences. In this calibration, people have constant relative risk aversion (CRRA) utility and a Prelec (1998) probability weighting function, and they can allocate their

portfolios across a positively skewed individual stock (portfolio), a negatively skewed mutual fund, and a risk-free asset. Our calibration is generally similar to Polkovnichenko (2005); details are provided in Online Appendix A. Figure 2 shows the optimal fraction of equity allocated to the individual stock, for different levels of probability weighting – denoted *Inverse-S* – and utility curvature. The fraction of equity allocated to the individual stock is strongly increasing in probability weighting. In contrast, the allocation is relatively insensitive to utility curvature. Thus, our simple calibrated portfolio choice model is consistent with prior theoretical papers that predict people with high *Inverse-S* will hold underdiversified portfolios with high positive skewness.

[Insert Figure 2 here]

Furthermore, Epstein and Zin (1990), Segal and Spivak (1990), and Chapman and Polkovnichenko (2011) show that probability weighting creates first-order risk aversion – that is, the investor does not become locally risk neutral as the size of a potential investment becomes small. This increased risk aversion reduces demand for equity securities.

1.3. The Elicitation Procedure

Empirically estimating individual-level measures of probability weighting is complicated because preferences are determined by the product of two (usually non-linear) functions: probability weighting and utility. Throughout the paper we will use the less conventional term “utility curvature” to refer to aversion to risk caused by utility curvature, and not the more frequently used term “risk aversion,” as technically, with probability weighting the curvature of the utility function alone does not fully describe risk aversion. Risk averse behavior is the outcome of a combination of utility curvature and probability weighting.

Thus, the challenge is to separate the effects of probability weighting from utility function curvature. For elicitation questions with modest rewards this issue is trivial if the subject integrates

outcomes with existing wealth as in expected utility theory or RDU, because the subject's utility function is effectively linear for modest rewards and its curvature can be ignored.⁶ Under behavioral theories that involve narrow framing, however, the subject evaluates decisions in isolation and utility function curvature can affect even small stake gambles. Hence, separating utility curvature and probability weighting is not trivial when assuming narrow framing. Prior studies address this issue using two different approaches. First, parametric methods that assume a specific functional form and then estimate probability weighting and utility curvature parameters (e.g., Tanaka, Camerer, and Nguyen, 2010; Erner, Klos, and Langer, 2013). The disadvantages of this approach are the need to commit to a specific functional form and the estimation error in the individual level parameter estimates. Second, non-parametric methods that do not assume a functional form, but which require chaining, so that the choices offered to a subject depend upon her prior choices (e.g., Wakker and Deneffe, 1996; Abdellaoui, 2000; van de Kuilen and Wakker, 2011). The disadvantage of this approach is that, as Abdellaoui (2000, pg. 1511) notes "...error propagation in the trade-off method can produce 'noisy' probability weighting functions" (e.g., a response error in the first question affects the choices offered in all subsequent questions).

We use a non-parametric approach and limit the need for chaining. Our survey questions are adapted from Wakker and Deneffe (1996) and Abdellaoui (2000), albeit with some modifications that reduce error propagation and the time required to complete the questions, due to the constraints of a general population survey (rather than a classroom experiment).⁷ We

⁶ This does not mean that individuals are effectively risk neutral for small gambles, however, as probability weighting alone generates first-order risk aversion. But deviations from risk neutrality are caused by probability weighting and not utility function curvature (e.g., see Yaari, 1987; Segal and Spivak, 1990).

⁷ We first piloted four different designs of the elicitation method in a sample of 207 ALP respondents, comparing the method of Abdellaoui (2000) with the midweight method of van de Kuilen and Wakker (2011), while using two different question presentation formats (choice lists and multiple pairwise choices). For our main survey, we chose the question format that the respondents found clear, minimized mistakes, and led to lower average response times. Online Appendix B provides further details of the elicitation method. We do not include the pilot sample responses in our empirical tests and the subjects for the pilot were not included in the sample for the main survey.

designed and fielded a customized module in the American Life Panel (ALP) survey that presents subjects with 10 multi-round questions. The first four questions measure utility curvature and the remaining six measure probability weighting. Each question asks subjects to choose between two options: A and B. There are three rounds per question, and based on each subject's choice in a round, one option in the subsequent round is changed to become either more or less attractive. As a starting point for each question, we use the answer of a risk-neutral expected utility maximizer. Hence the choices offered to our subjects are determined only by their prior answers within the rounds of a single question, rather than across different sets of questions.

Figure 3 shows the first round of the first question, which is intended to measure utility curvature. Option A offers a 33% chance of winning \$12 and a 67% chance of winning \$3, while Option B initially offers a 33% chance of winning \$18 and a 67% chance of winning \$0. Accordingly, both options have an expected value of \$6 and offer the same chance of winning the larger payoff (33%), but Option B is riskier (Option B is a mean-preserving spread of Option A). If the subject selects the safer Option A, then Option B is made more attractive by increasing the winning amount to \$21. If, instead, the subject chooses Option B, then Option B is made less attractive by decreasing the winning amount to \$16. This process continues for three rounds, until the subject's indifference point is closely approximated. For each question, the subject is then presented with a fourth choice used only to evaluate consistency with prior choices.

Panel A of Table 1 shows the structure of the four sets of questions designed to measure utility curvature. In all four questions, the probability of winning the large prize is fixed at 33% for both Option A and B, but the potential winning amount increases with each question. Thus, the effect of probability weighting is largely cancelled out in the comparison between Options A and

B, as the probability of winning is the same. We ask four sets of questions instead of one to obtain a more accurate measure of utility curvature and minimize the effect of measurement error.

[Insert Table 1 and Figure 3 here]

We next present each subject with six questions designed to measure probability weighting. The goal is to elicit the certainty equivalent of Option A, which is a risky choice with two possible outcomes. Figure 4 depicts the first round of one of the questions: Option A offers a fixed large payoff of \$42 with probability $p = 5\%$ and a small payoff of \$6 with probability 95%, while Option B offers a sure amount of \$8 in the first round. If the subject chooses risky Option A, then in the second round the sure amount for Option B is increased to \$9. If the subject instead chooses Option B, then in the second round the sure amount is reduced to \$7. This process is repeated for three rounds until the certainty equivalent for Option A is closely approximated, as illustrated by the decision tree in Figure 5. In the remaining five sets of probability weighting questions, the probabilities, p , of winning the large prize in Option A are 12%, 25%, 50%, 75%, and 88%. For each question, we elicit the certainty equivalent and compare it to the expected value of the risky gamble so as to estimate the percentage risk premium.⁸ Panel B of Table 1 shows the structure of the six sets of probability weighting questions.

[Insert Figures 4 and 5 here]

The subjects in our survey module could win real rewards based on their choices. This is important, as prior studies show that real rewards produce more reliable estimates of preferences (Smith, 1976). At the beginning of the survey, the subjects were told that one of their choices would be randomly selected and played for real money. We paid a total of \$16,020 in real

⁸ For the four sets of risk aversion questions the certainty equivalent is not known, as the respondent compares two lotteries. For these questions, we define the % risk premium as the percentage difference between the respondent's elicited indifference value and the indifference value of a risk-neutral decision maker.

incentives to 2,072 of the 2,703 eligible subjects who completed the survey. The American Life Panel (ALP) was responsible for determining and making the incentive payments, and subjects in the ALP regularly participate in and receive payment from the ALP surveys. The involvement of the ALP should minimize the subjects' potential concerns about the credibility of the incentives.

An advantage of our experimental survey approach is that we can explicitly state the probabilities, ensuring the subjects know the precise probabilities of all outcomes. This allows us to measure *preferences* towards probabilities rather than *beliefs* about probabilities; in contrast, it is extremely difficult to disentangle preferences and beliefs in purely observational studies. For instance, the popularity of actuarially unfair extended warranties could result from either probability weighting or overestimation of the probability of failure (e.g., Abito and Salant, 2017).

1.4. *The Probability Weighting Measure*

Using the six indifference values elicited from the probability weighting questions described above, we create a probability weighting measure for each individual. First, we convert the indifference values into percentage premiums relative to the expected value of the risky gamble (Option A). For example, consider the 5% probability weighting question. Suppose a subject is indifferent between Option A [5%, \$42; 95%, \$6] and Option B [100%, \$8.25]. The expected value of Option A is \$7.80, implying a percentage risk premium for question $PW_{5\%}$ of: $(7.80 - 8.25)/7.80 = -5.8\%$. In this case, the premium is negative as the subject overweights the low probability of winning a large prize and demands a certainty equivalent greater than the expected value of the risky gamble.

The risk premiums are summarized in the final column of Panel B in Table 1. On average, for high probabilities, people demand large positive risk premiums. For small probabilities (5% and 12%), however, people are willing to pay more than the expected value to own the lottery.

This pattern is consistent with overweighting of small probabilities. In contrast, this pattern is inconsistent with any model of expected utility, including models that incorporate skewness preferences (see Quiggin, 1993).

Using these premiums, we create our non-parametric probability weighting variable, *Inverse-S*, as follows:

$$\textit{Inverse-S} = (PW_{88\%} + PW_{75\%} + PW_{50\%}) - (PW_{25\%} + PW_{12\%} + PW_{5\%}). \quad (4)$$

In the experimental literature, individuals switch from overweighting to underweighting probabilities in the range between 25% and 50% (where the probability weighting function crosses the diagonal in Figure 1).⁹ Thus this measure is simply the premiums in the overweighting range less the premiums in the underweighting range. Higher values indicate a more pronounced *Inverse-S* shape for the probability weighting function.

This measure is parsimonious and it allows us to avoid assuming a specific functional form for the probability weighting function. If individuals frame narrowly and utility function curvature affects the responses, taking the difference between the percentage premiums reduces the influence of curvature, because curvature affects all premiums similarly and is thus largely differenced out of the measure. The cost of the tradeoff we have made in our survey design – limiting chaining to avoid measurement error – is that it is theoretically possible for utility curvature to influence our non-parametric measure of probability weighting. In practice, however, this possibility does not appear to significantly affect the measure. As the summary statistics in the next section show, *Inverse-S* does not have a high correlation with our measure of utility curvature ($r = 0.09$), and our empirical results are theoretically inconsistent with *Inverse-S* measuring utility curvature.

⁹ Note that a positive risk premium for the 25% question does not necessarily imply that the 25% probability is underweighted. Instead, the effects of utility curvature may fully offset the effects of probability weighting.

Nevertheless, to ensure that *Inverse-S* does not inadvertently measure utility curvature, in robustness tests we also jointly estimate utility curvature and probability weighting parameters using parametric models.

Specifically, we estimate utility curvature using CRRA utility and the probability weighting function proposed by Prelec (1998, Eq. 3.1).¹⁰ The Prelec function has clear axiomatic foundations and is well suited for very small and very large probabilities. The curve features a fixed intersection point at $p = 1/e = 0.37$, which is consistent with experimental findings. As an additional robustness test, we also estimate the parameter of the salience function proposed by Bordalo, Gennaioli, and Shleifer (2012, p. 1255). The salience function provides an intuitive psychological foundation for why probability weighting occurs and its predictions are consistent with empirically observed patterns. Online Appendix C provides details about the estimation of both functions.

2. Data and Variables

2.1. Data Sources: American Life Panel Survey and CRSP

We fielded our survey module in the RAND American Life Panel¹¹ from June 20 to July 19, 2017. The ALP includes several thousand households that regularly answer Internet surveys, and households lacking Internet access at the recruiting stage are provided with a laptop and wireless service to limit selection bias. To ensure that the sample is representative of the U.S. population, we use survey weights provided by the ALP for all analyses and summary statistics

¹⁰ Tversky and Kahneman (1992) propose an alternative functional form, but their specification suffers from two significant drawbacks. First, for certain parameter values, the decision weights can be negative, resulting in a preference for first-order stochastically dominated gambles (Ingersoll, 2008). Second, their specification generates an artificial negative correlation between the curvature parameter of the utility function and the probability weighting parameter (Fehr-Duda and Epper, 2012). Nevertheless, in unreported results, we estimate the Tversky and Kahneman (1992) probability weighting parameter and find that it produces empirical results similar to those of the Prelec (1998) parameter. Results are available upon request.

¹¹ For further information about the ALP see Online Appendix D and <https://www.rand.org/labor/alp.html>.

reported in this paper. In addition to the probability weighting variables, our module also collects information on portfolio choice and some control variables. Other controls, such as demographic and economic characteristics, are available from earlier ALP survey modules. The ALP invited 3,397 panel members and closed out the survey when 2,701 completed the survey, resulting in a completion rate of 79.5%.

Respondents who indicated that they held individual stocks were asked to list the names (or tickers) of their five largest holdings. We match these names or tickers by hand to the CRSP daily stock return database,¹² and we then construct various measures of stock characteristics using daily return data from July 1, 2016 to June 30, 2017. We select this specific period since our survey was fielded from June 20 to July 19, 2017. Table 2 provides summary statistics of the key variables used in our study (Appendix Table A1 defines the variables).

[Insert Table 2 here]

2.2. *Dependent Variables*

Fraction Allocated to Individual Stocks Conditional is the fraction of the respondent's total equity portfolio invested in individual stocks, conditional upon non-zero equity ownership. Among stock market participants, the average fraction allocated to individual stocks was 45%. Calvet, Campbell, and Sodini (2007, 2009) show that this fraction is a good proxy for portfolio underdiversification. In addition, for the subsample of individual stock owners, we observe the number of individual companies that they own. We find that, conditional on owning individual stocks, half of the respondents held shares in only one or two individual companies, which is

¹² In our tests, we use only U.S. based common stocks. We are unable to match 12.1% of the reported holdings because the holding was a foreign or private company, or because the reported name was ambiguous or unmatchable.

consistent with the fraction of the equity portfolio allocated to individual stocks being a reasonable proxy for underdiversification.

As an alternative measure of portfolio underdiversification, we also calculate the *Relative Sharpe Ratio Loss (RSRL)* of each respondent (following Campbell, Calvet, and Sodini, 2007, Eq. 7). We assume that the investor's mutual fund holdings are in a market index fund (beta of one and no idiosyncratic risk) and calculate the *RSRL* as follows:

$$RSRL_i = 1 - \frac{\mu_i/\sigma_i}{\mu_M/\sigma_M} = 1 - \frac{\beta_i \cdot \sigma_M}{\sigma_i}, \quad (6)$$

where μ_i (μ_m) is the risk premium of the investor's portfolio (market portfolio), σ_i (σ_M) is the standard deviation of the investor's portfolio (market portfolio), and β_i is the beta of the investor's entire portfolio. One caveat is that we do not know the exact amount invested in each individual stock, we know only the total amount invested in individual stocks. Hence, we assume that the investor holds an equally weighted portfolio of individual stocks. The investor's *RSRL* will equal zero if he holds a fully diversified portfolio while larger values indicate underdiversification.

We also generate several stock level measures of (expected) skewness using the CRSP return data. *Total Skewness* is the skewness of daily stock returns. Following Kumar (2009), *Idiosyncratic Skewness* is the skewness of the residuals from a two-factor model that includes the market risk premium, *RMRF*, and its square, *RMRF*². We include the square of the market risk premium to remove loading on systematic skewness. *Max. One-Day Return* is the maximum one-day return over the period, which Bali, Cakici, and Whitelaw (2011) argue is a good proxy for investors' beliefs about lottery-like payoffs. *Idiosyncratic σ* is the annualized standard deviation of the residuals from a Fama and French (2015) five-factor model. *Stock β* is the average market beta of the investor's stock holdings. For respondents who own multiple stocks, we first average across stocks for that person and then average across respondents, for the summary statistics

reported in Table 2. We perform the tests at the individual stock level and at the portfolio level using the equal-weighted daily returns of the investor's stockholdings.

The final three dependent variables in Table 2 are summarized for the entire sample, including subjects who do not own equities. *Mutual Funds Only* is an indicator variable equal to one for the 8.3% of the respondents who own only equity mutual funds and no individual stocks. *Individual Stocks Only* is an indicator variable equal to one for the 7.0% of the sample whose equity ownership consists exclusively of individual company stocks and no equity mutual funds. *Both Mutual Funds and Individual Stocks* is an indicator variable equal to one for the 8.6% of the sample who own both equity mutual funds and individual stocks.¹³

2.3. Control Variables

All of the empirical tests control for demographic and economic characteristics including age, sex, race, ethnicity, marital status, number of household members, education, employment status, family income, and financial wealth. Including these variables controls for the potential confounding effects they might have on household portfolio choice.

Our ALP survey module also included additional questions to measure utility curvature, numeracy, financial literacy, trust, and optimism.¹⁴ These variables mitigate against the potential omitted variable bias that could occur from factors that are conceptually similar to probability weighting. For example, utility curvature could be highly correlated with probability weighting.

¹³ Our sample had a lower equity participation rate than that reported in some other studies as we exclude equity ownership in 401(k) and other tax deferred plans. Such equity holdings may not reflect active choices by the respondent, due to the U.S. Department of Labor's acceptance of target date funds as an investment default. This permits employees to be defaulted into holding equities instead of making active choices. For more on target date funds and 401(k) plan investment options, see Mitchell and Utkus (2012). Further, employee ownership of publicly-traded companies occurs primarily through tax deferred plans such as 401(k) and employee stock ownership plans (see Curcuro, Heaton, Lucas, and Moore, 2010).

¹⁴ Online Appendix D provides the exact wording of the questions to measure numeracy, financial literacy, trust, and optimism.

Thus, in the regressions, we control utility curvature to ensure that our probability weighting variable captures a component of preferences that is distinct from utility curvature. Our measure of utility curvature is the average of the risk premiums from the four utility curvature questions summarized in Panel A of Table 1.

Optimism could influence the overweighting of small probabilities (i.e., optimists may assume that small probability lotteries always resolve in their favor). Accordingly, we follow Puri and Robinson (2007) and include a question assessing individuals' subjective life expectancies; this permits us to measure optimism by comparing subjective and objective life expectancies (where the latter are derived from age/sex population mortality tables). We also control for financial literacy which prior studies show has a strong association with financial decisions (e.g., Lusardi and Mitchell, 2007, 2014; van Rooij, Lusardi, and Alessie, 2011). To ensure that overweighting of small probabilities is not simply a proxy for low financial literacy, our survey module also includes the "Big Three" financial literacy questions developed by Lusardi and Mitchell (2007) for the Health and Retirement Study (HRS). Our index of financial literacy is the number of correct responses to these questions; on average, respondents answered slightly more than two questions correctly. The module also includes three questions to assess numeracy based on the HRS and the English Longitudinal Study of Ageing, along with the trust question from the World Values Survey, as Guiso, Sapienza, and Zingales (2008) report a relation between trust and portfolio choice.

2.4. Probability Weighting

Panel B of Table 1 summarizes the responses to the six probability weighting questions from the ALP survey module. On average, subjects are risk-seeking for low probability questions with $p = 0.05$ and $p = 0.12$; indeed, consistent with the overweighting of small probabilities, the

average risk premiums are negative (-7.1% and -2.3%, respectively). For these questions, the required positive risk premium due to utility curvature is more than offset by the risk-seeking due to probability weighting. For the $p = 0.25$ question, the average risk premium is 4.6%. At larger probabilities, $p = 0.5, 0.75$ and 0.88 , the average risk premiums increase to 15.1%, 22.8%, and 28.2%, respectively. Overall, the pattern in the average risk premiums is consistent with *Inverse-S*-shaped probability weighting: overweighting of small probabilities and underweighting of high probabilities.

Our probability weighting variable, *Inverse-S*, is positive for 81% of the respondents, indicating an inverse-S shaped probability weighting function,¹⁵ which is consistent with the results from laboratory experiments using students (e.g., Abdellaoui, 2000; Bruhin, Fehr-Duda, and Epper, 2010). Panel A of Table 3 summarizes the non-parametric probability weighting measure, *Inverse-S*. Consistent with inverse-S shaped probability weighting in the general population, on average the sum of the risk premiums for the three high probability questions exceeds the sum of the risk premiums for the three low probability questions by 71 percentage points. Panel A also shows there is substantial heterogeneity in probability weighting, which has important implications for the finance literature as it may help explain the observed large heterogeneity in portfolio allocations.

[Insert Table 3 here]

Panel B of Table 3 shows the pairwise correlations between *Inverse-S* and education, utility curvature, numeracy, financial literacy, optimism, and trust. Although not the main focus of our paper, these correlations help illustrate the underlying distribution of our *Inverse-S* measure. The correlation between utility curvature and *Inverse-S* is low and positive ($r = 0.092$),

¹⁵ Similarly, when we fit the Prelec (1998) weighting function jointly with a CRRA power utility function using all ten questions, the majority (73%) of the respondents exhibit an inverse-S shaped function (see Online Appendix C).

with utility curvature explaining less than 1% (R^2) of the variation in *Inverse-S*. To place this small correlation in perspective, the average correlation among the risk premiums of the four utility curvature questions in Panel A of Table 1 is $r = 0.70$ (demonstrating strong internal consistency). Accordingly, *Inverse-S* and utility curvature appear to be separate components of preferences.

The correlations in Table 3 also provide evidence on the relation between *Inverse-S* and proxies for intelligence. Panel B shows that *Inverse-S* is positively correlated with education, numeracy, and financial literacy, although the magnitudes are small. Thus there is no evidence that probability weighting is greater for individuals who are less intelligent or less educated; this is consistent with Booi, van Praag, and van de Kuilen (2010) who also find no relation between education and probability weighting in the general population.

Optimism could potentially lead to overweighting the probability of winning the lotteries. However, this would decrease the risk premiums for all questions instead of generating a pattern of risk-seeking for low probabilities and risk aversion for high probabilities. Because we construct *Inverse-S* as the difference between risk premiums, any influence from optimism should be approximately differenced out. Indeed, we do not find a significant correlation between *Inverse-S* and optimism.

3. Probability Weighting and Household Portfolio Choice

Next, we test the relation between probability weighting and household portfolio choice decisions. For ease of interpretation, we standardize the *Inverse-S* variable so it has a mean of zero and a standard deviation of one. Following Dimmock, Kouwenberg, Mitchell, and Peijnenburg (2016), all specifications include controls for age, age squared, education, log(family income), log(financial wealth), sex, white, Hispanic, log(number of household members), employment

status, and dummies for variables for which we imputed missing values. Our baseline specifications also include controls for utility curvature, numeracy, financial literacy, optimism, and trust.¹⁶ For all specifications, we report t -statistics calculated using robust standard errors.

3.1. *Probability Weighting and Equity Portfolio Underdiversification*

Table 4 shows the results of Tobit regressions in which the dependent variable is a measure of portfolio underdiversification. In Panel A, the dependent variable is the fraction of the subject's equity holdings (individual stocks plus stock mutual funds) invested in individual stocks (see Polkovnichenko, 2005; Calvet, Campbell, and Sodini, 2007). In Panel B, the dependent variable is the *Relative Sharpe Ratio Loss* variable of Calvet, Campbell, and Sodini (2007). In both panels, the sample includes only those subjects with non-zero equity holdings. Column (1) includes no additional variables; column (2) adds the economic and demographic controls; column (3) adds our measure of risk aversion; and Column (4) adds controls for numeracy, financial literacy, optimism, and trust.

[Insert Table 4 here]

Theoretical models predict that probability weighting will make underdiversified portfolios more attractive due to their positive skewness (e.g., Shefrin and Statman, 2000; Polkovnichenko, 2005; Barberis and Huang, 2008; Jin and Zhou, 2008). As predicted by theory, Panel A shows a significant and positive relation between *Inverse-S* and the fraction of equity holdings allocated to individual stocks. The results are similar across all four columns and there is little change as additional control variables are included. The coefficient reported in column (4) implies that a one

¹⁶ Dimmock, Kouwenberg, Mitchell, and Peijnenburg (2016) find that ambiguity aversion relates to portfolio choices of households. In our elicitation questions, all probabilities are known, so there is no ambiguity. For some respondents, we have their measure of ambiguity aversion. The correlation between *Inverse-S* and ambiguity aversion is 0.05 and our results do not change when we include this variable as control (results available upon request).

standard deviation increase in *Inverse-S* results in a 12.7 percentage point increase in the fraction of the portfolio allocated to individual stocks (a 28.2% increase relative to the baseline allocation of 45.0 percentage points).

Panel B also show a significant positive relation between probability weighting and portfolio underdiversification. In all columns, higher *Inverse-S* is associated with higher *Relative Sharpe Ratio Loss*. Thus individuals who overweight small probability tail events hold portfolios with lower Sharpe ratios than could have been obtained with similar levels of systematic risk. The coefficient reported in column (4) implies that a one standard deviation increase in *Inverse-S* results in a 4.3% lower Sharpe ratio, relative to the market index. To interpret the economic magnitude of these results, we use the return loss measure of Calvet, Campbell, and Sodini (2007, Eq. 11). Our results imply that for a one-standard deviation increase in *Inverse-S* the average (median) stockholder loses \$2,504 (\$351) per year.¹⁷

The results in Panel B are generally similar to those in Panel A, though the sample size is smaller because some respondents do not provide stock identifiers or the identifiers cannot be matched to specific individual stocks.¹⁸ Given the similarity of the results, and because the two proxies for underdiversification have a correlation of 0.9, in the remainder of the paper we report results only for the *Fraction of Equity in Individual Stocks*.

¹⁷ The dollar return loss is the additional expected dollar return an investor could have received given her overall level of risk. It is calculated by fixing the investor's overall portfolio risk, but replacing the (uncompensated) idiosyncratic risk with (compensated) systematic risk.

¹⁸ In particular, 40 respondents did not report the name or ticker of any of their holdings and 56 respondents gave names or tickers that could not be matched as they reported holdings that were not domestic common stocks or could not be matched to a single security.

3.2. *Alternative Measures of Probability Weighting Preferences*

Our main analyses use a parsimonious non-parametric measure for the *Inverse-S* parameter.¹⁹ As a robustness test, we estimate two alternative versions of the baseline specification in which we replace *Inverse-S* with a parametrically-estimated probability weighting measure. First, we use the functional form proposed by Prelec (1998, Eq. 3.1) because it has clear axiomatic foundations and is widely used in the decision sciences literature. In this specification, we use a parametric measure of risk aversion that was jointly estimated along with the probability weighting parameter. Second, we use the functional form proposed by Bordalo, Gennaioli, and Shleifer (2012, p. 1255) in their model of salience because it provides a clear and intuitive explanation for why probability weighting occurs. The correlations between the non-parametric *Inverse-S* measure and the alternative measures by Prelec (1998) and Bordalo, Gennaioli, and Shleifer (2012) are 0.75 and 0.78, respectively.

The results are presented in Panel A of Table 5. Both of the parametric measures are defined so that higher values indicate greater probability weighting. Consistent with earlier tables, we standardize these variables to have a mean of zero and standard deviation of one. For both parametric measures of probability weighting, the results are similar to those in the main specification. Hence, our results are robust to using this alternative method to separate probability weighting from utility curvature.

¹⁹ To demonstrate that the results are not driven by outliers, Appendix Table A2 reports results from specifications that use three alternative measures of *Inverse-S*: a rank transformation for which zero indicates the lowest value of *Inverse-S* and one the highest, an indicator variable equal to one if the subject's *Inverse-S* value is above the median, and an indicator variable equal to one if the subject's *Inverse-S* value is above 25% (in which case we can reasonably assume the person has *Inverse-S* shaped probability weighting preferences). In all three specifications there is a positive and significant relation between *Inverse-S* and the fraction of equity allocated to individual stocks.

3.3. *Measurement Error in Preference Elicitation*

A large literature, beginning with Harless and Camerer (1994) and Hey and Orme (1994), shows that subjects often give inconsistent responses to questions designed to elicit preferences. If such errors are pure noise, this will reduce the power of tests but not introduce bias. If, however, errors in elicited preference are correlated with errors in actual decisions (e.g., holding an undiversified portfolio), this could potentially affect inferences. Our ALP module includes several features that allow us to empirically address this issue. First, we included the check questions described earlier to test the internal consistency of subjects' choices. Second, the ALP module recorded the amount of time subjects spent on each question, which allows us to identify subjects who answered the elicitation questions unusually quickly. Accordingly, Panel B of Table 5 reports results for two restricted subsamples, which exclude subjects who made more than three mistakes on the check questions or who spent less than 1.5 minute on the probability weighting questions. The results are similar to those in the full sample for both of the restricted samples, suggesting that our main results are not driven by measurement error in elicited preferences or by individuals who failed to understand the elicitation questions.

[Insert Table 5 here]

4. Probability Weighting and Individual Stock Characteristics

Probability weighting has implications not just for the choice between mutual funds and individual stocks, but also for the *type* of individual stocks an investor chooses. Investors who overweight the probabilities of tail events should select individual stocks with high positive skewness, but they will not exhibit a preference for high systematic risk (e.g., see Barberis and Huang, 2008; Boyer, Mitton, and Vorkink, 2010). Investing in a positively skewed stock is

appealing because the investor has a chance, albeit a small one, of becoming rich if that company becomes the “next Apple.”

Our survey module asked subjects who owned individual stocks to list the names (or tickers) of their five largest individual stock holdings. The five largest holdings encompass the entire portfolio of most individual stockholders in the sample; about half hold only one or two stocks, and 75% hold five or fewer. As described in Section 2.2., we match these stocks to the CRSP daily stock return database and construct various stock characteristics that measure skewness: *Total Skewness*, *Idiosyncratic Skewness*, *Max. One-Day Return*, and *Idiosyncratic σ* . Note that we include *Idiosyncratic σ* because prior papers argue it is a good proxy for expected positive skewness (e.g., Boyer, Mitton, and Vorkink, 2010), and not because probability weighting implies a preference for idiosyncratic risk itself. We also calculate the market beta, *Stock β* , as a measure of systematic risk.

Table 6 shows regression results for the five dependent variables described above. The key independent variable is *Inverse-S*, the sample includes only subjects with individual stockholdings, and all models include the full set of controls. In Panel A, the unit of observation is an individual stockholding (e.g., there are three observations for an investor who holds three stocks) and standard errors are clustered by individual. In this panel, the focus is on the specific type of stocks selected. In Panel B, the unit of observation is the investor’s entire equity portfolio, and the dependent variables are characteristics calculated from the returns of an equally-weighted portfolio of the investor’s stockholdings combined with equity mutual fund holdings. In this panel, the focus is on the investor’s overall equity portfolio and the characteristics account for portfolio level diversification effects.

[Insert Table 6 here]

Columns (1) to (3) show that *Inverse-S* is significantly and positively related to *Total Skewness*, *Idiosyncratic Skewness*, and *Max. One-Day Return*. Investors with higher probability weighting choose stocks that have high expected positive skewness. Column (4) of Panel A shows that *Inverse-S* has a positive and significant (at the 10% level) relation with idiosyncratic risk (a proxy for expected skewness), however this relation is not significant in the portfolio level results shown in Panel B.

Column (5) shows that the relation between *Inverse-S* and systematic risk, measured by *Stock β* , is neither statistically nor economically significant. Thus, the overall pattern of results in Table 6 is that investors with high *Inverse-S* choose to hold portfolios with high expected positive skewness but not higher systematic risk. Importantly, this pattern is precisely what is implied by probability weighting. It is not, however, an obvious implication of alternative explanations. For example, if *Inverse-S* inadvertently measured risk-seeking preferences, it would imply higher positive skewness and higher systematic risk.

These results also relate to two streams of the literature that argue probability weighting explains observed behaviors in financial markets. First, although our data do not allow us to directly test the relation between probability weighting and asset pricing, the results in this section are consistent with studies of positive skewness and asset pricing. For instance, Boyer, Mitton, and Vorkink (2010), Bali, Cakici, and Whitelaw (2011), Conrad, Dittmar, and Ghysels (2013), Conrad, Kapadia, and Xing (2014), and Barberis, Mukherjee, and Wang (2016) show that stocks with positive expected skewness have abnormally low returns.²⁰ Barberis and Huang (2008) argue that probability weighting can cause positively skewed securities to have low returns. Our results support the findings of these asset pricing studies, by providing direct evidence that

²⁰ Relatedly, Wang (2017) shows undervaluation of securities with negative expected skewness.

individuals who overweight small probabilities exhibit a preference for positively skewed securities. Second, the results in this section are consistent with several papers that argue financial institutions design products that cater to investor's probability weighting preferences (Henderson and Pearson, 2011; Li, Subrahmanyam, and Yang, 2018). These products are popular despite having large negative abnormal returns.

5. Preference, Probability Unsophistication, or Financial Knowledge?

Thus far, we have interpreted probability weighting as a component of *preferences* rather than a misunderstanding of the underlying probabilities. In other words, we posit that subjects (approximately) understand the underlying probabilities but, when maximizing their preference functions, they use weighted rather than actual probabilities. In this section, we consider two closely-related alternative explanations based on miscalibration of the underlying probabilities rather than preferences towards probabilities: (1) Probability unsophistication – that some individuals have difficulty with probabilistic reasoning, and this difficulty affects both their elicited *Inverse-S* values and their portfolio choices; and (2) Limited financial knowledge – that, for some reason, *Inverse-S* is correlated with a lack of financial knowledge. For these alternative explanations, probability weighting and underdiversification represent clear mistakes. That is, if high *Inverse-S* respondents were educated about probabilistic reasoning or financial markets, they would likely make different choices. With preferences, on the other hand, such interventions would not result in different choices.²¹

²¹ Note that even if probability weighting does reflect preferences, it can still be considered a mistake, as it constitutes a violation of the independence axiom. Nevertheless this it is a fundamentally different type of mistake and one that is more difficult to change.

5.1. *Preference or Probability Unsophistication*

The first alternative explanation is based on probability unsophistication. For example, subjects with limited quantitative reasoning skills may have difficulty evaluating questions involving probabilities, and such cognitive limitations could also cause investment errors. This explanation appears unlikely based on the summary statistics presented in Panel B of Table 3, which show that *Inverse-S* has a small but significantly *positive* correlation with education, numeracy, and financial literacy. Nevertheless, we perform additional tests using two restricted samples. In column (1) of Table 7, the sample includes only subjects with a college degree (undergraduate or graduate). In column (2), the sample includes only subjects who correctly answer all three of the numeracy questions. In both restricted samples, and consistent with the full sample results, *Inverse-S* has a significant positive relation with portfolio underdiversification. These results suggest that *Inverse-S* does not reflect poor quantitative reasoning.

[Insert Table 7 here]

5.2. *Preference or Limited Financial Knowledge*

The second alternative explanation is that *Inverse-S* may be correlated with a lack of financial knowledge. For example, some subjects may simply be unaware of the benefits of diversification or fail to understand financial risks, resulting in investment errors,²² and this investment behavior may be correlated with *Inverse-S*. A priori, however, it is not obvious why a lack of financial knowledge would be correlated with *Inverse-S*. Indeed, a key advantage of eliciting probability weighting preferences using lotteries instead of natural events is that we can clearly and unambiguously define the relevant probabilities, limiting the scope for beliefs to affect subjects' responses (for further discussion see Barberis, 2013b, p. 614).

²² For instance, von Gaudecker (2015) finds that underdiversification is related to low financial literacy.

We test this alternative using two restricted samples. In column (3) of Table 7, the sample includes only subjects who correctly answer all three financial literacy questions. In column (4), the sample includes only subjects who correctly answer the question “Please tell us whether this statement is true or false. ‘Buying a stock mutual fund usually provides a safer return than a single company stock.’” In the latter sample, subjects correctly state that a mutual fund is usually safer than an individual stock. The results in these restricted samples are generally similar to those in the full sample. *Inverse-S* is positively associated with portfolio underdiversification, even for investors who understand the risks associated with underdiversification. These subjects do not choose individual stocks because they misunderstand the risks; rather, they choose to hold individual stocks despite knowing they are riskier than mutual funds.

Overall, this section suggests that the relation between *Inverse-S* and underdiversification is due to preferences towards probabilities, rather than probabilistic unsophistication or a lack of knowledge about the risks of underdiversification. This finding contributes to the discussion about whether probability weighting constitutes a “mistake,” in the sense that if people were told that their behavior violated an axiom, they would be expected to reverse their choices (e.g., Fehr-Duda and Epper, 2012; Barberis, 2013a; Barberis, 2013b). Our results are consistent with early experimental studies finding that people are unwilling to change choices violating the independence axiom even after the axiom is explained to them (see MacCrimmon, 1968; Slovic and Tversky, 1974).

6. Probability Weighting and Participation in Mutual Funds and Individual Stocks

Next, we broaden the analysis to consider non-participation in equity markets, as well as the choice between individual stocks versus stock mutual funds by those who do participate. For these tests, the theoretical predictions are less clear than for the tests discussed above. If an

investor's choice set includes only the risk-free asset and a mutual fund (or market index), prior studies show that probability weighting can cause non-participation due to first-order risk aversion²³ (e.g., Chapman and Polkovnichenko, 2011; De Giorgi and Legg, 2012; He, Kouwenberg, and Zhou, 2018). When an individual stock is added to the choice set, however, theory is not well developed and the predictions are less clear. On the one hand, probability weighting implies first-order risk aversion, which makes any equity investment, including individual stocks, less attractive. On the other hand, probability weighting implies a preference for positive skewness, which makes individual stocks more attractive. Thus, probability weighting can result in non-participation or portfolio underdiversification, depending on a subject's beliefs about the risks and skewness of individual stocks. Hence the net effect of probability weighting is an empirical question.

Table 8 shows the results of multinomial logit models in which the dependent variable, *Portfolio Choice*, takes one of four values: *Non-Participation*, *Mutual Funds Only*, *Individual Stocks Only*, and *Both Mutual Funds and Individual Stocks*. Classic financial theory predicts that individuals should (1) participate in the equity market and (2) hold a well-diversified portfolio (e.g., the market portfolio). In other words, the rational benchmark is to invest in mutual funds; both non-participation and underdiversification are possible behavioral deviations from rationality due to probability weighting. Accordingly, *Mutual Funds Only* serves as a natural basis of comparison and we use it as the excluded or reference category in the multinomial model.

[Insert Table 8 here]

²³ Epstein and Zin (1990), Segal and Spivak (1990), and Chapman and Polkovnichenko (2011) show that the probability weighting in rank-dependent utility implies first-order risk aversion (i.e., that the investor does not become locally risk neutral as the size of a potential investment becomes small).

The results show that subjects with higher *Inverse-S* are more likely to choose either non-participation or individual stock ownership, and they are thus less likely to own only equity mutual funds. The economic magnitudes implied by the coefficient estimates are large. For instance, the coefficient in column (1) implies that a one standard deviation increase in *Inverse-S* raises the probability of choosing *Non-Participation* instead of *Mutual Funds Only* by one-third ($e^{0.290} = 1.34$). Likewise, a one standard deviation increase in *Inverse-S* raises the probability of choosing *Individual Stocks Only* instead of *Mutual Funds Only* by 39.8%, and choosing *Both Mutual Funds and Individual Stocks* instead of *Mutual Fund Only* by 31.1%.

The interpretation of the multinomial logit results comes with a caveat, however, as theoretically whether high *Inverse-S* results in non-participation or underdiversification depends on the subject's beliefs about expected equity returns, risk, and individual stock skewness (e.g., He, Kouwenberg, and Zhou, 2018). Unfortunately, we lack data on beliefs about return distributions, and so we cannot disentangle why some high *Inverse-S* subjects choose not to participate in the stock market while others buy positively-skewed individual stocks.

We emphasize that the pattern of results in Table 8 is broadly consistent with the theoretical predictions of probability weighting, while it is inconsistent with most alternative interpretations of our measure. For example, if *Inverse-S* inadvertently measured utility curvature (risk aversion), it would be positively related to non-participation but negatively related to underdiversification. Alternatively, if *Inverse-S* inadvertently measured optimism, it would be negatively related to non-participation. Instead, however, *Inverse-S* is positively related to non-participation.

7. Conclusion

This paper is the first to provide *direct* evidence linking probability weighting and portfolio underdiversification. We measure probability weighting using real incentives in a survey module

fielded in a large and nationally representative sample of the U.S. population. Using our new *Inverse-S* metric, we demonstrate that most individuals exhibit probability weighting – they overweight low probability tail events – though there is also substantial heterogeneity. Next we evaluate how probability weighting preferences relate to household portfolios, and find that higher probability weighting is associated with underdiversification. Among investors who own individual stocks, we also show that probability weighting is associated with higher positive stock skewness. Further, we find evidence consistent with probability weighting being a component of preferences, rather than the result of probability unsophistication or lack of financial knowledge.

The implied economic magnitudes of our results are large; a one-standard deviation higher *Inverse-S* implies a cost to the average (median) stockholder of \$2,504 (\$351) per year. Furthermore, probability weighting increases the dispersion of portfolio returns, pushing people to either not participate or hold positively skewed portfolios. This will lead to increased heterogeneity in realized investment returns, which potentially exacerbates wealth inequality (e.g., Bach, Calvet, and Sodini, 2017; Lusardi, Michaud, and Mitchell, 2017; Fagereng et al. 2018).

References

- Abdellaoui, Mohammed. 2000. Parameter-Free Elicitation of Utility and Probability Weighting Functions. *Management Science* 46, 1497-1512.
- Abito, Jose Miguel, and Yuval Salant. 2017. The Effect of Product Misperception on Economic Outcomes: Evidence from the Extended Warranty Market. Working paper, University of Pennsylvania.
- Albuquerque, Rui. 2012. Skewness in Stock Returns: Reconciling the Evidence on Firm Versus Aggregate Returns. *Review of Financial Studies* 25, 1630-1673.
- Allais, Maurice. 1953. Le Comportement de L'Homme Rationnel Devant le Risque: Critique des Postulats et Axiomes de L'Ecole Américaine. *Econometrica* 21, 269-290.
- Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2017. Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy, Working paper, Stockholm School of Economics.
- Bali, Turan G., Nusret Cakici, and Robert Whitelaw. 2011. Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns. *Journal of Financial Economics* 99, 427-446.
- Barberis, Nicholas C. 2013a. Thirty Years of Prospect Theory in Economics: A Review and Assessment. *Journal of Economic Perspectives* 27, 173-196.
- Barberis, Nicholas C. 2013b. The Psychology of Tail Events: Progress and Challenges. *American Economic Review (Papers & Proceedings)* 103, 611-616.
- Barberis, Nicholas C., and Ming Huang. 2008. Stocks as Lotteries: The Implications of Probability Weighting for Security Prices. *American Economic Review* 98, 2066-2100.
- Barberis, Nicholas C., Abhiroop Mukherjee, and Baolin Wang. 2016. Prospect Theory and Stock Returns: An Empirical Test. *Review of Financial Studies* 29, 3068-3107.
- Blume, Marshall E., and Irwin Friend. 1975. The Asset Structure of Individual Portfolios and Some Implications for Utility Functions. *Journal of Finance* 30, 585-603.
- Booij, Adam S., Bernard M.S. van Praag, and Gijs van de Kuilen. 2010. A Parametric Analysis of Prospect Theory's Functionals for the General Population. *Theory and Decision* 68, 115-148.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. Saliency Theory of Choice under Risk. *Quarterly Journal of Economics* 127, 1243-1285.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. Saliency and Asset Prices. *American Economic Review (Papers & Proceedings)* 103, 623-628
- Boyer, Brian, Todd Mitton, and Keith Vorkink. 2010. Expected Idiosyncratic Skewness. *Review of Financial Studies* 23, 169-202.
- Boyer, Brian, and Keith Vorkink. 2014. Stock Options as Lotteries. *Journal of Finance* 69, 1485-1527.
- Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper. 2010. Risk and Rationality: Uncovering Heterogeneity in Probability Distortion. *Econometrica* 78, 1375-1412.

- Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. 2007. Down or Out: Assessing the Welfare Costs of Household Investment Mistakes. *Journal of Political Economy* 115, 707-474.
- Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. 2009. Measuring the Financial Sophistication of Households. *American Economic Review (Papers & Proceedings)* 99, 393-398
- Camerer, Colin. 1995. Individual Decision Making, In: J. Kagel and A. Roth, eds., *Handbook of Experimental Economics* Princeton: Princeton University Press: 587-703.
- Chapman, David A., and Valery Polkovnichenko. 2011. Risk Attitudes Towards Small and Large Bets in the Presence of Background Risk. *Review of Finance* 15, 909-927.
- Conrad, Jennifer, Robert F. Dittmar, and Eric Ghysels. 2013. Ex Ante Skewness and Expected Stock Returns. *Journal of Finance* 68, 85-124.
- Conrad, Jennifer, Nishad Kapadia, and Yuhang Xing. 2014. Death and Jackpots: Why Do Individual Investors Hold Overpriced Stocks? *Journal of Financial Economics* 113, 455-475.
- Curcuro, Stephanie, John Heaton, Deborah Lucas, and Damien Moore. 2010. Heterogeneity and Portfolio Choice: Theory and Evidence. In: Y. Ait-Sahalia and L.P. Hansen, editors, *Handbook of Financial Econometrics: Tools and Techniques*, Elsevier, 337-382.
- De Giorgi, Enrico G., and Shane Legg. 2012. Dynamic Portfolio Choice and Asset Pricing with Narrow Framing and Probability Weighting. *Journal of Economic Dynamics and Controls* 36, 951-972.
- Dimmock, Stephen G., Roy Kouwenberg, Olivia S. Mitchell, and Kim Peijnenburg. 2016. Ambiguity Aversion and Household Portfolio Choice Puzzles: Empirical Evidence. *Journal of Financial Economics* 119, 559-577.
- Epstein, Larry G., and Stanley E. Zin. 1990. 'First-Order' Risk Aversion and the Equity Premium Puzzle. *Journal of Monetary Economics* 26, 387-407.
- Erner, Carsten, Alexander Klos, and Thomas Langer. 2013. Can Prospect Theory be Used to Predict an Investor's Willingness to Pay? *Journal of Banking and Finance* 37, 1960-1973.
- Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2018. Heterogeneity and Persistence in Returns to Wealth. NBER working paper 22822.
- Fama, Eugene F., and Kenneth R. French. 2015. A Five-Factor Asset Pricing Model. *Journal of Financial Economics* 116, 1-22.
- Fehr-Duda, Helga, and Thomas Epper. 2012. Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences. *Annual Review of Economics* 4, 567-593.
- Goetzmann, William N., and Alok Kumar. 2008. Equity Portfolio Diversification. *Review of Finance* 12, 433-463.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2008. Trusting the Stock Market. *Journal of Finance* 63, 2557-2600.

- Harless, David W., and Colin F. Camerer. 1994. The Predictive Utility of Generalized Expected Utility Theories. *Econometrica* 62, 1251-1289.
- He, Xue Dong, Roy Kouwenberg, and Xun Yu Zhou. 2018. Inverse S-Shaped Probability Weighting and Its Impact on Investment. *Mathematical Control & Related Fields*, forthcoming.
- Henderson, Brian J., and Neal D. Pearson. 2011. The Dark Side of Financial Innovation: A Case Study of the Pricing of a Retail Financial Product. *Journal of Financial Economics* 100, 227-247.
- Hey, John D., and Chris Orme. 1994. Investigating Generalizations of Expected Utility Theory Using Experimental Data. *Econometrica* 62, 1291-1326.
- Ingersoll, Jonathan. 2008. Non-Monotonicity of the Tversky-Kahneman Probability-Weighting Function: A Cautionary Note. *European Financial Management* 14, 385-390.
- Jin, Hanqing, and Xun Yu Zhou. 2008. Behavioral Portfolio Selection in Continuous Time. *Mathematical Finance* 18, 385-426.
- Kahneman, Daniel, and Amos Tversky. 1979. Prospect Theory: An Analysis of Decisions Under Risk. *Econometrica* 47, 263-291.
- Kelly, Morgan. 1995. All Their Eggs in One Basket: Portfolio Diversification of US Households. *Journal of Economic Behavior and Organization* 27, 87-96.
- Kumar, Alok. 2009. Who Gambles in the Stock Market? *Journal of Finance* 64, 1889-1933.
- Li, Xindan, Avanidhar Subrahmanyam, and Xuewei Yang. 2018. Can Financial Innovation Succeed by Catering to Behavioral Preferences? Evidence from a Callable Options Market. *Journal of Financial Economics* 128, 38-65.
- Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell. 2017. Optimal Financial Knowledge and Wealth Inequality. *Journal of Political Economy* 125, 431-477.
- Lusardi, Annamaria, and Olivia S. Mitchell. 2007. Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth. *Journal of Monetary Economics* 54, 205-224.
- Lusardi, Annamaria, and Olivia S. Mitchell. 2014. The Economic Importance of Financial Literacy: Theory and Evidence. *Journal of Economic Literature* 52, 5-44.
- MacCrimmon, Kenneth. 1968. Descriptive and Normative Implications of the Decision Theory Postulates. *Risk and Uncertainty* 3, 3-32.
- Mitchell, Olivia S., and Steve P. Utkus. 2012. Target-Date Funds in 401(k) Retirement Plans. NBER working paper 17911.
- Mitton, Todd, and Keith Vorkink. 2007. Equilibrium Underdiversification and the Preference for Skewness. *Review of Financial Studies* 20, 1255-1288.
- Polkovnichenko, Valery. 2005. Household Portfolio Diversification: A Case for Rank-Dependent Preferences. *Review of Financial Studies* 18, 1467-1502.

- Prelec, Drazen. 1998. The Probability Weighting Function. *Econometrica* 66, 497-527.
- Puri, Manju, and David T. Robinson. 2007. Optimism and Economic Choice. *Journal of Financial Economics* 86, 71-99.
- Quiggin, John. 1982. A Theory of Anticipated Utility. *Journal of Economic Behavior and Organization* 3, 323-343.
- Quiggin, John. 1993. *Generalized Expected Utility Theory: The Rank Dependent Model*. Kluwer: Norwell, MA, USA.
- Rieger, Marc O. 2012. Why Do Investors Buy Bad Financial Products? Probability Misestimation and Preferences in Financial Investment Decision. *Journal of Behavioral Finance* 13, 108-118.
- Segal, Uzi, and Avia Spivak. 1990. First Order versus Second Order Risk Aversion. *Journal of Economic Theory* 51, 111-125.
- Shefrin, Hersh, and Meir Statman. 2000. Behavioral Portfolio Theory. *Journal of Financial and Quantitative Analysis* 35, 127-151.
- Slovic, Paul, and Amos Tversky. 1974. Who Accepts Savage's Axiom? *Systems Research and Behavioral Science* 19, 368-373.
- Starmer, Chris. 2000. Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. *Journal of Economic Literature* 38, 332-382.
- Smith, Vernon L. 1976. Experimental Economics: Induced Value Theory. *American Economic Review* 66, 274-279.
- Tanaka, Tomomi, Colin F. Camerer, and Quang Nguyen. 2010. Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam. *American Economic Review* 100, 557-571.
- Tversky, Amos, and Daniel Kahneman. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty* 5, 297-323.
- van de Kuilen, Gijs and Peter P. Wakker. 2011. The Midweight Method to Measure Attitudes Toward Risk and Ambiguity. *Management Science* 57, 582-599.
- van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie. 2011. Financial Literacy and Stock Market Participation. *Journal of Financial Economics* 101, 449-472.
- von Gaudecker, Hans-Martin. 2015. How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice? *Journal of Finance* 70, 489-507.
- Yaari, Menahem E. 1987. The Dual Theory of Choice under Risk. *Econometrica* 55, 95-115.
- Wakker, Peter P., and Daniel Deneffe. 1996. Eliciting von Neumann-Morgenstern Utilities When Probabilities are Distorted or Unknown. *Management Science* 42, 1131-1150.
- Wang, Baolian. 2017. Probability Weighting and Asset Prices: Evidence from Mergers and Acquisitions. Working paper, Fordham University.