A New Way to Measure Consumer Risk Preferences

Daniel G. Goldstein
London Business School

Eric J. Johnson
Graduate School of Business
Columbia University

William F. Sharpe
Graduate School of Business
Stanford University

Address all correspondence to:
Professor Daniel Goldstein
London Business School
Regent's Park
London, NW1 4SA
UK
Phone: +44 (0)20 7706 6796
Fax: +44 (0) 20 7724 1145
Email: dgoldstein@london.edu

A New Way to Measure Consumer Risk Preferences

Consumer choice occurs over multiple products and services, each comprising multiple risks. In this paper, we present a new market research technique to measure consumers' preferences over large spaces of risks. We first describe the method, present its psychological and analytical motivation, and then report the results of empirical tests of reliability and validity, both within testing sessions and across the span of one year. The method is used to estimate the coefficient of relative risk aversion and the loss aversion parameter for a sample of adults saving for retirement. The new technique passes tests of reliability and validation and captures individual differences based on age and income. It also identifies two sub-populations, one best fit by a more classical model of risk preference, and the other by a behavioral model which incorporates loss aversion.

Keywords: marketing research tools, consumer behavior, decision-making, parameter estimation, measurement, segmentation, risk, utility, uncertainty

Introduction

Consumer choice occurs over multiple products and services, each having multiple associated risks (Bauer 1960; Bettman 1973). The numerous outcomes and probabilities facing consumers may be unknown (Kahn and Sarin 1988), or as we shall explore, be known but be too numerous and interdependent for the unaided mind to process. Product and service risks cause consumers anxiety (Locander and Hermann 1979; Sheth and Venkatesan 1968; Roselius 1971) as they worry about matters such as automotive breakdown, computer failure, drug toxicity, and fund underperformance. Marketing managers have similar worries from the sell side, along with additional concerns about predicting risk attitudes in order to design promotions, loyalty programs (Biyalogorsky and Gerstner 2004; Kivetz 2003), and sales contests (Kalra and Shi 2001; Gaba and Kalra 1999).

Despite the importance of risk to marketing, the measurement of risk preferences has lagged behind the swift progress made in the measurement of riskless, attribute-level preferences. Specifically, conjoint analysis has rocketed from its humble roots in mathematical psychology to become the marketing's chief methodological export (Luce and Tukey 1964; Green, Krieger, and Wind 2004). The lack of a risk measurement method in the marketer's toolbox might limit the impact that marketing research can have on insurance, financial services, and medicine, in which risk is an inextricable component. This paper presents a new market research technique for measuring consumer preferences over complex multiple-outcome risks.

How do people think about risk and return? Since the Enlightenment, theories of risk preference have been based on choices between simple gambles. In the last century, choices among gambles have been used to show violations of Utility Theory (Allais 1953), and motivate subsequent alternatives, such as Subjective Expected Utility Theory (Savage 1954), Rank Dependent Utility Theory (Quiggin 1982), Cumulative Prospect Theory (Tversky and Kahneman 1992), Reference-Dependent Subjective Expected Utility Theory (Sugden 2003), among many others (e.g., Birnbaum and McIntosh 1996; Lopes and Oden 1999). In recent years, alternative models to the alternative models have been motivated by choices between gambles as well (Wu 1994; Loomes and Sugden 1998; Birnbaum, Patton, and Lott 1999).

Nearly all behavioral studies of risk look at what we will call ***simple prospects***, that is, single prospects (a probability $\underline{p}$ of outcome $\underline{x}$), and binary prospects, (a probability $\underline{p}$ of outcome $\underline{x}$ and probability $\underline{q}$ of outcome $\underline{y}$). Studies of this kind are so common that researchers have found it worthwhile to model the heuristics people use specifically when choosing between simple prospects: editing (Kahneman and Tversky 1979), cancellation, combination, simplification, elimination, segregation, and a variant of the Take The Best Heuristic (Gigerenzer and Goldstein

1996), to name a few (Brandstaetter, Gigerenzer, and Hertwig in press; Kahneman and Tversky 1979; Wu 1994)[1].

With the most notable exception being the estimation of consumer utility functions (Hauser and Urban 1979), why has consumer research had so little to say on risk preference? One possibility might be that simple prospects, while analytically tractable, are not realistic descriptions of the risks consumers regularly face. First, real product and service risks have multiple outcomes, not just two: An investment can return any percentage of its principal, and an insurance policy can be worth any percentage of its cost. Second, unlike simple prospects, consumer choice often occurs over far more than two alternatives. A consumer shopping for funds with Fidelity has over 4,500 investment products from which to choose. Each fund has a continuous distribution of outcomes, and each can be combined with other funds to create more portfolios than could be enumerated in a lifetime. As Lopes (1987) puts it, simple prospects "occur most frequently in the context of formal gambling and psychology experiments." Measuring risk in the domain of multiple-alternative, multiple-outcome prospects seems warranted.

In this paper, we focus on investing for retirement, one of the largest decisions many people ever make, and certainly one of the consequential consumer decisions (Kahn and Luce 2003) for which researchers have been hearing the call. Most employees in the United States, upon starting a new job, are taken down the hall to the human resources department where they are shown a list of investment products, and asked to allocate 100 percentage points of their retirement contribution between them. Many employees spend less than an hour deciding how to allocate assets for retirement, which is surprisingly little time considering that as many as 90% or more will never change their initial choice (Benartzi and Thaler 2001; Samuelson and Zeckhauser 1988) and that the decision could impact their well-being for one third of their lives or more.

Economic theory (Merton 1969, 1971; Samuelson 1969; Hakansson 1970), investment advice (Perold and Sharpe 1988), and the marketing collateral of mutual fund vendors often recommend that people divide their contributions between a risky and risk-free (or very low risk) asset and perform periodic rebalancing to fixed proportions by dollar value. If returns are uncorrelated over time such an investment strategy yields a roughly lognormal distribution of long-term returns.[2] If consumers follow this strategy their wealth at retirement will, effectively, be drawn from a lognormal distribution. However, this begs an important question: how many

[1] We focus on multiple outcomes on a single dimension; there is a substantial literature on multi-attribute prospects, which we will not address in this article.

[2] If returns are independent, the ending value for such a strategy will converge to lognormality as the number of periods increases. If r is the rate of return, the ending value is the product of (1+r), its log is the sum of the logs of (1+r) and the central limit theorem holds that the sum of uncorrelated random variables converges to normality.

investors would state a preference for a lognormal distribution of wealth? And is a lognormal distribution the most suitable alternative given most people's stated preferences? Some investors may desire more downside protection, or more upside potential. If it turns out to be the case that a substantial number of people are loss averse (Benartzi and Thaler 1995; Kahneman and Tversky 1979) and prefer downside-protected, thus non-lognormal distributions, it would mean that many investors are not getting the kinds of outcomes they say they want. It would also suggest market opportunities realized only in part by current theory and practice. We present an inquiry into consumers' preferences for investment outcomes to see how many lay investors' preferences align with what common investment advice recommends.

The aim of this paper is to present a new market research technique for studying preferences over multiple-outcome risks. We first describe the method, present its psychological and analytical motivations, and then report the results of empirical tests of its reliability and validity both within testing sessions and across the span of one year. Empirically, we use this method to estimate the coefficient of relative risk aversion and the loss aversion parameter for a sample of working adults who have been saving for retirement for 5 to 30 years. To foreshadow our results, the method passes tests of reliability and validation and captures individual differences based on age and income. It also identifies two sub-populations, one best fit by the classical economic theory of risk preference, and the other by a behavioral model incorporating loss aversion. We conclude by discussing how the new methodology can impact research on risk and consumer decision-making.

The Distribution Builder

Psychological Rationale

The Distribution Builder (DB) elicits preferences over multiple-outcome risks without requiring participants to engage in any numerical estimation or computation. Though the method can be applied over categorical outcomes (such as whether or not a call will go through on a wireless network), we focus here on a single continuous criterion. Customer expectation distributions (Rust et al. 1999), being distributions, can easily be captured in the DB framework. Other examples in the consumer domain include the distribution of: delivery times for an online purchase, defects in a shipment, fraudulent transactions in auctions, times until product failure, response times under service contracts, queuing times at health care facilities, and the example domain of this article, wealth resulting from investment products.

The Distribution Builder displays risks as discrete probability distributions, and assesses preferences by asking the consumer to build a desirable distribution, subject to constraints. DB aims to help consumers construct preferences in eliminating, by design, a number of pitfalls that are endemic to probability elicitation and utility assessment. The method is grounded in three distinctions in the psychology of information representation: frequency versus numerical representations of probability, experienced versus stated representations of probability, and, integrated versus isolated judgments.

Frequency versus numerical representations of probability. The various ways in which probabilities are displayed (as fractions, decimals, or relative frequencies), though mathematically equivalent, are psychologically different in the computations they facilitate (Feynman 1967; Kahneman, Slovic, and Tversky 1982). For example, changing numerical probabilities to natural frequencies (for example, changing .001 to “1 in 1000”) has been shown to cause a shift from incorrect answers to normatively correct ones in certain base-rate neglect problems. For a review of research on frequencies, see the work of Gigerenzer and colleagues (1994; 2002). Frequency representations have been found to facilitate probabilistic reasoning in a number of domains including: making medical diagnoses (Hoffrage and Gigerenzer 1998), understanding medical test results (Hoffrage, Kurzenhäuser, and Gigerenzer 2005), presenting DNA evidence in court (Hoffrage et al. 2000), and communicating weather reports (Gigerenzer et al. 2005). The Distribution Builder makes use of frequency representations exclusively.

Experienced versus stated representations of probability. Psychologically, there are important differences in the perception of probabilities that are stated as summary statistics—such as reading the number .25—and those that are experienced—such as seeing 25 of 100 coin flips turn up heads (Christensen-Szalanski and Bushyhead 1981; Manis et al. 1980; Sedlmeier 1997). Koehler (1996) gives a review of studies showing that that trial-by-trial learning of base rate information increases normative responses in base rate neglect tasks. For example, Gigerenzer, Hell and Blank’s (1988) investigation allowed people to experience random sampling directly in Kahneman & Tversky’s (1973) classic Engineer-Lawyer problem and as a result, base rate neglect diminished. A psychological basis for learning probabilities through experience is given by Hasher and Zacks (1984), who demonstrate the surprisingly acute ability of people to monitor natural frequencies, and by Spellman and colleagues (Spellman 1996; Holyoak and Spellman 1993; Spellman 1993) who argue that repeated trials are automatically encoded for use by the implicit learning system. In DB, trial-by-trial random drawing is not simply claimed, it is visually observable.

Integrated versus isolated representations. Decisions made in isolation can result in inconsistent choices. Examples of this include demonstrations that two sequential decisions (for instance, risks or gambles), can result in different choices than a single integrated choice (Thaler and Johnson 1990), and that choices made about isolated risks can lead to undesirable outcomes when those risks are aggregated (Kahneman and Lovallo 1993; Tversky and Shafir 1992). In the

DB framework, interrelations between sets of gambles are explicitly displayed: one risk cannot be explored without observing its effect on the entire distribution of risks. Unlike in some gamble choice tasks, potential gains on the upside must be paid with potential losses on the down side. Furthermore, the DB representation makes it impossible to submit probabilities that sum to less or more than one (so-called subadditive or superadditive probabilities, see Rottenstreich and Tversky 1997; Tversky and Fox 1995) , a common problem in utility assessment (Baron 1988).

The Distribution Builder Interface

To illustrate how these psychological affordances are built into the interface and to explain the relationship between responses on the DB and preferences, we start by exploring the interface, shown in Figure 1.

[Insert Figure 1 about here]

Design consideration: frequency representations of probability. Outcomes are expressed by the rows of the vertical axis, and probabilities by the number of markers placed on each row. In Figure 1, the 41 rows represent outcomes labeled 0% to 200%. As an example, these outcomes might represent percentages of pre-retirement income earned annually in retirement. The lowest row has 3 markers stacked at 65%, representing a .03 probability of obtaining the outcome labeled 65%. By clicking and dragging a marker, the user can move that one probability marker and all those to its right. Thus, users can quickly and easily alter the probabilities assigned to various outcomes, and an n-outcome distribution can be constructed in n movements of the mouse. To prevent biasing participants towards a particular distribution, at the start of the session all markers are off the board, in a separate row. In a typical study, participants are told to construct a gamble that they themselves would like to play. Participants are then instructed to move the markers into a desirable pattern (in actuality, a multiple outcome distribution), such that they are comfortable knowing that any one of the markers could turn out to represent them. They are told that the computer has randomly selected one of the markers, but that it is impossible to know in advance which one.

Design considerations: experienced representation of probabilities. When participants are finished with the distribution they have constructed, they click the submit button, and ***all but one*** of the 100 markers disappear one at a time, leaving the randomly selected marker showing. The payoff of the prospect is equal to the row on which the one remaining marker is standing. In this way, participants visually experience random drawing from a probability distribution.

Design consideration: integrated representations via cost functions. While the interface is straightforward to use, it can represent complex prospects. With 100 markers and 40 rows, DB can represent over 10^{34} ways of assigning probabilities to outcomes[3]. However, to provide a representation that can be modeled as a utility function, additional constraints must be provided in the form of a cost function. To explain how, we define some terminology. A Distribution Builder consists of N markers that may be placed on M ***outcome rows*** each with a corresponding ***outcome row value***. The first (or lowest or bottom) outcome row corresponds to the lowest outcome row value. In Figure 1, M equals 41 outcome rows with outcome row values ranging from 0% to 200% in increments of 5 percentage points. In referring to the ith marker, we mean the ith marker from the bottom of the distribution. Markers are counted from the lowest outcome row to the highest, and from left to right within a row. Each time a marker is moved, the counting begins anew from the bottom. Let $\boldsymbol{o_i}$ be the outcome row value (or outcome for short) of the ith marker. Outcomes are always assigned to satisfy $\boldsymbol{o}_1 \leq \boldsymbol{o}_2 \leq \ldots \leq \boldsymbol{o_N}$, that is, they are indexed in increasing order. Let $\boldsymbol{s_i}$ be the cost associated with the ith marker, and the cost of a given distribution be $\sum_{\boldsymbol{i}=1}^{\boldsymbol{N}} \boldsymbol{o_i s_i}$. In order to limit the space of possible distributions by cost, DB can be configured such that the participant is not allowed to submit a distribution until the cost of the distribution equals some predetermined budget B, that is, until the following cost equation is satisfied:

(1) $$\sum_{\boldsymbol{i}=1}^{\boldsymbol{N}} \boldsymbol{o_i s_i} = \boldsymbol{B}$$

As the user experiments with different distributions, the ***cost meter*** seen in Figure 1 displays the value of the overall cost of the distribution graphically and numerically. If the cost of the distribution is too low, the meter turns blue; too high and the meter turns red. When the cost is within an acceptable margin around B, the meter turns green, and only then may the distribution be submitted.

Two examples illustrate cost functions in practice. Suppose a researcher wishes to limit participants to choose among only those distributions with an expected value B. This would be achieved by setting all the $\boldsymbol{s_i}$ equal to $\boldsymbol{p_i}$, where $\boldsymbol{p_i}$ is the probability of the ith marker being drawn. In our examples $\boldsymbol{p_i}$ is always 1/100, making the cost of the distribution equivalent to its expected value. Given a distribution that satisfies this particular cost constraint, if one of the

[3] There are $\binom{\boldsymbol{N}+\boldsymbol{r}-1}{\boldsymbol{r}-1}$ ways to partition N objects into r categories.

markers is moved up one outcome row, then somewhere in the distribution, one marker must be moved down one row. The distribution must balance around its expected value, so that which is added to the upside must be taken from the downside. Interestingly, this is not the case with all cost vectors. Suppose the goal is to have a cost function that allows a given amount of downside risk to be rewarded with an even greater amount of upside gain (as in financial markets). With a cost vector such as $1 > \boldsymbol{s}_1 > \boldsymbol{s}_2 > \ldots > \boldsymbol{s}_{\boldsymbol{N}} > 0$ the lower-valued outcomes are weighted more heavily than the higher valued ones, and have a larger impact on the cost of the distribution. Moving one marker down one row may enable the user to move several markers up several rows. As a consequence, with such a cost vector, the expected value of the distribution can actually ***increase*** as more downside risk is taken on, because it must be compensated with proportionally more of the lesser-weighted upside gain in order to satisfy the cost constraint in Equation 1. We now illustrate use of the DB in an applied consumer domain.

Using the Distribution Builder to Assess Preferences For Risks Associated With Investment Products

Our empirical example focuses on investing for retirement, a domain of great consequence, as there are nearly 2 trillion dollars invested in 401(k) retirement savings accounts in the United States. While we will concentrate on the retirement investment decision in this paper, it should be noted that the DB is a broadly applicable method for eliciting distributions and has no special relationship to finance.

Figure 1 shows a DB configured to measure preferences for wealth in retirement. The outcome row values, ranging from 0 to 200%, represent annual income in retirement expressed as a percentage of pre-retirement income. For someone with a pre-retirement income of $100,000, the 75% outcome row would correspond to an annual income of $75,000 per year. Participants in our experiment were asked to create a distribution they would like to have apply to their own retirement income. Distributions were constrained by cost according to the function $\sum_{\boldsymbol{i}=1}^{\boldsymbol{N}} \boldsymbol{o}_{\boldsymbol{i}} \boldsymbol{s}_{\boldsymbol{i}} = \boldsymbol{B}$ where the budget B is the exactly enough money needed to retire at 75% of pre-retirement income if only a risk-free investment is used, and $\boldsymbol{s}_{\boldsymbol{i}}$ is the Arrow-Debreu state price of the i^{th} marker (Arrow 1964; Debreu 1959). The underlying assumption is that there are N states of the world and markets that explicitly or implicitly allow the purchase of a ***state claim*** on each state. Thus a claim on state i will pay $1 if and only if state i occurs. The market price of such a claim is s_i. In order to minimize the cost of any distribution chosen by the user of the DB, markers are assigned to the N states with the lowest outcomes obtained in the cheapest states. Recall that outcomes are indexed in increasing (or non-decreasing) order. We thus index state prices in decreasing (or non-increasing) order. The least-cost way to obtain a distribution of outcomes is thus $\sum_{i=1}^{N} o_i s_i$, which is the cost associated with the distribution and shown as a percent of the budget (B) on the cost meter.[4]

A key attribute of this procedure is that state prices are decreasing in value so that taking on downside risk allows proportionally more upside gain (as in real markets). In particular, using the method of state prices, the budget meter will reflect the cost of the distribution according to a least-cost investment strategy (Dybvig 1988). While any decreasing set of state prices could provide useful information about participants' preferences, it is desirable to choose trade-offs that are similar to those in actual capital markets. To do so, the state prices were based on a ten-year investment horizon with real bond returns equal to 2% per year and real stock returns lognormally distributed with an annual mean of 8% per year and a standard deviation of 18% per year.[5]

Given the computed set of state prices, the participant's budget is set to the exact amount required to finance a certain outcome of 75%. That is, the least-variance distribution that will satisfy the cost constraint is the one in which all 100 markers are on the 75% row, shown in Figure 2. Distributions comprising a wider range of outcomes will have higher expected values, as is the case with the distribution in Figure 1.

[Insert Figure 2 About Here]

Using distributions to identify characteristics of consumers' utility functions

Distribution Builder enables the collection of rich data on subjective utility functions from distributions instead of from choices between simple prospects. In order to do this we assume that a person building a distribution is maximizing the function $\boldsymbol{p} \bullet \boldsymbol{u}(\boldsymbol{o})$ subject to the cost constraint $\sum_{i=1}^{N} o_i s_i = B$. Maximizing the Lagrangian $\boldsymbol{p} \bullet \boldsymbol{u}(\boldsymbol{o}) - \boldsymbol{k}(\boldsymbol{o}' \bullet \boldsymbol{s} - \boldsymbol{B})$ results in the family of equations $p_i u'(o_i) = ks_i$ for each marker i, where k is a constant. If we let K be k divided by any of the identical p_i (in our example 1/100), this becomes:

(2) $$u'(o_i) = Ks_i$$

[4] For working papers describing early experiments using the DB, see Sharpe and colleagues (Sharpe, Goldstein, & Blythe 2000; Sharpe 2001).
[5] for details, see Sharpe (2001).

Thus, the slope of the utility function at point i differs from the ith state price by a constant, a result that will later allow us to estimate loss aversion.

The important result here is that, given assumptions, the assignment of markers to wealth levels provides information about the marginal utility of each wealth level for an individual, enabling the estimation of individual utility functions and parameter under various models of utility. As a practical consumer concern, we can test whether participants' utility is consistent with advice often given to retirement investors.

One common set of assumptions used in investment advice is that investors' risk preferences possess a property called Constant Relative Risk Aversion, or CRRA (Pratt 1964; Safra and Segal 1998; Barberis 2000; Arrow 1970). The three most commonly used utility functions in financial economic models are the quadratic, the constant absolute risk aversion and the constant relative risk aversion (CRRA). The quadratic function implies that an investor will put fewer dollars in riskless assets as he or she becomes wealthier. The constant absolute risk aversion function implies that an investor will put the same number of dollars in riskless assets as wealth increases. The constant relative risk aversion function implies that the proportion of wealth invested in riskless assets is invariant with respect to changes in initial wealth level. Of the three functions, CRRA provides implications closest to observed behavior on the part of most investors (Arrow 1970; Pratt 1964).

People with CRRA utility functions should prefer a constant mix of assets if returns are uncorrelated over time, resulting in lognormally-shaped distributions of terminal wealth. A typical CRRA utility function is the power utility function $\boldsymbol{u}(\boldsymbol{o}) = \frac{\boldsymbol{o}^{1-\alpha}}{1-\alpha}$. An investor with such a function will have a marginal utility of $\boldsymbol{u}'(\boldsymbol{o}) = \boldsymbol{o}^{-\alpha}$. Combining this with equation (2) and taking logarithms gives:

(3) $$\boldsymbol{ln}(s_i) = \boldsymbol{ln}(K) - \alpha\, \boldsymbol{ln}(o_i)$$

In the finance literature α is commonly referred to as the coefficient of relative risk aversion. If most people's risk preferences are well fit by CRRA, standard investment advice should help them select investment products. If not, then many people could be holding portfolios with risk characteristics that violate their expressed risk preferences. In the experiment which follows, we fit Equation (3) at an individual level. If the relationship is not linear (low R^2), it would suggest that investor preferences may be better described by a model other than that of CRRA. If the relationship is linear, we would expect estimates of the coefficient of relative risk aversion α to fall in the range commonly found by other means.

Experiment 1

The purpose of the first experiment is to estimate the parameters of risk aversion and loss aversion for preferences concerning wealth in retirement, and to obtain a descriptive overview of data gathered with DB.

Method and Participants. Experimentation was carried out via the Internet. Participants were 141 geographically diverse US residents and citizens who are members of Columbia University's Center for the Decision Sciences participant pool who were paid for their participation. Participants had an average age of 41 with a standard deviation of 8 years. All have been saving for retirement for at least 5 years, and were at least 5 years away from retirement. 76% were married, 12% single and 12% divorced or widowed. Median income was $50,000 with a standard deviation of about $33,000. Average net worth was estimated to be about $200,000 and average amount saved towards retirement was about $110,000.

Participants were instructed to think about income in retirement as a percentage of pre-retirement income. Based on common advice given to people saving for retirement, they were told that 75% of pre-retirement income was a typically recommended goal. In addition, participants were instructed to notice that the 75% row on the DB is highlighted to reflect it being a typical retirement income level goal. An extensive training session explained how to interact with the computer program, and participants created one practice distribution, which did not have any budget constraint (or visible cost meter), but did demonstrate the one-by-one random selection of markers once the distribution was submitted. After this, users were instructed on the role of the cost meter. They were told their task was to find a pattern of markers that they would have apply to their own income in retirement, and that uses between 99 and 100% of the budget. It was emphasized that it is very important to treat the task as if it concerned one's own income in retirement. Participants were strongly advised against taking a chance of going below 25% of pre-retirement income, and for this reason, as seen in Figures 1 and 2, the bottom rows of the DB are shaded. Participants created two distributions, one after the other, and then went on to answer a survey containing demographic questions.

In addition, because we are interested in how well the Distribution Builder does in relation to other techniques, we included a psychological risk profile: the participant's score on the gambling and investment subscales of the Weber-Blais-Betz Domain Specific Risk Attitude Scale (Weber, Blais, and Betz 2002). In addition, we contrasted our method with an industrial risk profile quiz, a variant of an actual risk tolerance self-assessment questionnaire used by one of the world's largest providers of retirement investment products. The text of these items is found in the appendix. In addition, participants completed two crucial validation tasks, the outcome preference task and the gamble choice task, which are described fully in Experiment 2.

Results

Distributions. Figure 3 shows the aggregate distribution based on the number of markers placed at each wealth level averaged across all participants and both distributions. The first distributions have a mean expectation of 97.5% of pre-retirement income and a mean standard deviation of 24.5%; for second distributions these values were 92.5% and 20.3%. Note that the composite distribution is roughly lognormal in shape, but exhibits a peak at the reference point at 75%, congruent with loss-averse preferences, as will be discussed in the next section.

[Insert Figure 3 about here]

Model fit. We fit a power function $\boldsymbol{u}(\boldsymbol{o}) = \frac{\boldsymbol{o}^{1-\alpha}}{1-\alpha}$, which is a CRRA utility function, to each participant's first and second distributions. Since the state prices (s_i values) were given and the participants chose a distribution of outcome (o_i) values, we treated the logarithm of the former as the independent variable and the logarithm of the latter as the dependent variable in our regressions. Re-arranging equation (3) shows that the slope in such a regression provides an estimate of $-\frac{1}{\alpha}$, which can be easily transformed into the coefficient of relative risk aversion α. Economic estimates of α have been found to range in the area from around 1 to 10 or higher (Campbell 1996; Brav, Constantinides, and Geczy 2002; Blake 1996; Mehra and Prescott 1985; Friend and Blume 1975). The first distributions had a median α of 4.41, the second distributions 6.27[6], well within the expected range. The median R^2 value across all first distributions was .92, and across second distributions was .89.

An alternative to the traditional CRRA model are loss averse utility functions such as that in Prospect Theory (Kahneman and Tversky 1979), which have received recent attention in the marketing domain (Novemsky and Kahneman 2005; Camerer 2005; Ariely, Huber, and Wertenbroch 2005). In loss-averse utility functions, losses have a higher impact than gains due to a loss aversion parameter λ, which applies on one side of a reference point perceived to be the border between gains and losses, and which has been empirically estimated to be around 2.25 (Tversky and Kahneman 1992). For

[6] We excluded cases where regression cannot compute a CRRA parameter because all 100 markers were placed on the same row.

instance, we could make a loss averse version of the CRRA power utility function by positing

(4) $\boldsymbol{u}'(\boldsymbol{o}) = \boldsymbol{K}\boldsymbol{o}^{-\alpha}$ in the domain of gains (o > the reference point), and

$\boldsymbol{u}'(\boldsymbol{o}) = \lambda\boldsymbol{K}\boldsymbol{o}^{-\alpha}$ in the domain of losses (o < the reference point).

The participants in this experiment had a clear reference point: the 75% row. Participants were instructed that 75% of pre-retirement income is a typical goal, and the 75% row was highlighted on the interface. Furthermore, 75% was also the risk-free alternative in the experiment: placing all 100 markers at 75% is the lowest variance distribution that satisfied the budget meter (see Figure 2). Because this served as a salient reference point, and because most participants placed outcomes above and below it, the loss aversion parameter λ could be estimated at an individual level. Recall that from Equation 2, the state price s_i of the i^{th} marker equals the slope of the utility function divided by a constant, that is, $\frac{\boldsymbol{u}'(\boldsymbol{o}_i)}{\boldsymbol{K}} = \boldsymbol{s}_i$. Owing to this result, λ can be estimated from the state prices of the markers immediately above and below the reference point:

(5) $$\frac{\boldsymbol{s}_i \text{ below reference point}}{\boldsymbol{s}_i \text{ above reference point}} = \frac{\frac{\boldsymbol{u}'(\boldsymbol{o})\text{ losses}}{\boldsymbol{K}}}{\frac{\boldsymbol{u}'(\boldsymbol{o})\text{ gains}}{\boldsymbol{K}}} = \frac{\lambda\boldsymbol{o}^{-\alpha}}{\boldsymbol{o}^{-\alpha}} = \lambda$$

Or, in more intuitive terms, the kink in the utility function parameterized by λ is the ratio of the slopes of the utility function as approached from the gain side and the loss side. In our data, across both distributions and all subjects, λ was 2.08, consistent with many other estimates (Tversky and Kahneman 1992). While this might be considered to indicate widespread occurrences of kinked utility curves, it is important to recall that participants were restricted to choices of outcomes at 41 levels, each of which differed by its successor and predecessor by 5% of income in retirement. Due to this granularity, λ values may slightly either overstate or understate the magnitudes of any kinks in participants' utility functions.

Individual differences: Loss aversion. Thus far, we have found α to be around 5 and λ around 2, reasonably in line with what would be expected. However, the question arises: do these aggregate estimates mask variation across people? In other words, are there subpopulations of people with different risk preferences? Are participants who are not well fit by CRRA loss averse, or do they deviate in some other way from the CRRA model?

Recall that CRRA-maximizing distributions plot as a straight line in log-log space, and that the median R^2 in the experiment centers around .9. If we perform a median split and categorize participants with R^2 >=.9 as CRRA and R^2 <.9 as non-CRRA, we can see whether the non-CRRA group appears loss averse. Across the participants in this investigation, looking just at the first distribution, the median λ for low R^2 participants is 2.22, and for high R^2 participants is 1.44. For the second distributions, the values are 3.97 and 1.54. Indeed, the people not described well by CRRA built markedly more loss-averse distributions.

The difference in the distributions between the two subpopulations of participants is also easily visually observed. Figure 4 shows the composite distributions of high and low R^2 participants[7]. The distribution of the high R^2 group is roughly lognormal in shape, skewed to the right. The low R^2 distribution deviates from this smooth form in a very specific way, consistent with high loss aversion. To understand why this shape suggests loss aversion, recall that the loss aversion parameter is estimated by comparing the state prices of the makers just below and just above the reference point. The markers at the reference point are excluded. When the number of markers at the reference point is large, the difference between the two state prices on either side of it becomes greater (because state prices are assigned in order of position in the distribution), and accordingly, so does the ratio. Note that this relationship between R^2 and loss aversion is not obvious. It is possible to create low R^2 distributions with low loss aversion, as is the case with a bimodal distribution, or a distribution with probability massed around a point other than the reference point.

[Insert Figure 4 about here]

Experiment 2

The purpose of Experiment 2 is to conduct tests of reliability and validation of DB measures. In order to do this, we had the same participants submit retest distributions one year after those submitted in Experiment 1.

Method and Participants. One year after submitting their first two distributions with DB, 85 participants from Experiment 1 and 70 participants from another DB experiment conducted at the same time as Experiment 1 completed a follow-up study, making a total of 155 year 2 respondents[8]. As in year 1, participants submitted two distributions concerning desired income in retirement. Next, they were presented a DB on which they were asked to play two small stakes gambles for a gain of up to $1.25 (USD) or a loss of up to $.75, with a risk free

---

[7] The average R^2 was computed per participant based on his or her two submitted distributions. Both distributions went into either the high or low R^2 composite histogram depending on this average.

[8] The unreported condition was exactly the same as the reported one except that the budget constraint was set to 60% instead of 75% of pre-retirement income. Jumping ahead a bit, there is no significant difference

alternative of $0 and outcome row increments of $.05. Outside of the axis labeling, the mechanics of the distribution builder were exactly the same as in the retirement scenarios. The stakes were real: money was added or subtracted to the participants' payment on the outcome of these gambles. After submitting the 2 retirement and 2 gamble distributions, participants engaged in two validation tasks.

Outcome preferences task. In the "outcome preference" task, participants were shown histograms representing distributions of returns resulting from portfolios invested 0%, 10%, 20%, 30%, 40%, 50%, or 60% in stock, and the rest in a risk-free asset[9] and asked which pattern of investment results they would like to have apply to their own retirement. Histograms of returns were chosen as a validating measure for four reasons: they have been shown to facilitate estimates of volatility (Ibrekk and Morgan 1987), they have been favored by people assessing risks (Thompson and Bloom 2000), they are commonly used in financial prospectuses and studies of financial risk perception (Siebenmorgen, Weber, and Weber 2000), and finally, unlike asset allocation tasks, they preclude application of the 1/N heuristic (Benartzi and Thaler 2001) which can bias responses. Because the histograms we present consist of a constant mix of stocks and a risk-free asset, they are lognormal in shape, and thus approximate the options available to people who follow the constant-mix investment advice.

Gamble choice task. To validate DB in comparison with the dominant form of risk assessment task, participants were presented 3 choices between sure amounts and gambles. The first item was "Which would you prefer if offered right now 1) $4.50 for certain 2) A 50% of getting $1 and a 50% chance of getting $15". The choices in the other items were "1) $1 for certain 2) A 10% chance of getting $12 and a 90% chance of getting nothing" and "1) $8 for certain 2) A 90% chance of getting $10 and a 10% chance of getting nothing." The number of risky choices (0-3) was computed for each participant.

Reliability and validation

Reliability. To assess the reliability of the risk measures collected with DB, we computed the reliability coefficient as the Pearson correlation between parameters generated by two distributions submitted one after the other, as shown in Table 1. Within the first session in year 1, the two measurements of the transformed CRRA model-based risk parameter $-1/\alpha$ have a Pearson correlation of .700 (Spearman .702). In the year 2 session, this correlation reached .803 (Spearman .793). Note that the distributions were completely reset between trials so that the high

---

in year 2 transformed alphas between participants from the 60% and 75% year 1 groups (t ratio -.4688, 144 DF).

[9] Graphs were made with the same assumptions about budget, holding period, and stock returns as used for the DB.

correlations are not due to mere persistence. To look at reliability over time, the average $-1/\alpha$ value from year 1 is correlated with its corresponding value in year 2, giving a Pearson correlation of .431 (Spearman .454). As Ghiselli, Campbell and Zedek (1986) point out, the between-year reliability measures $\boldsymbol{r}_{12(\boldsymbol{true})}$ must take into account the attenuation due to within-year reliabilities and be equivalent to $\frac{\boldsymbol{r}_{12}}{\sqrt{\boldsymbol{r}_{11}\boldsymbol{r}_{12}}}$ where $\boldsymbol{r}_{11}$ and $\boldsymbol{r}_{22}$ are the respective within year 1 and year 2 reliabilities. Substituting in the Pearson correlations as reliabilities here, $\boldsymbol{r}_{12(\boldsymbol{true})}$ is equal to .575. Reasoning that participants with little or no experience investing for retirement may provide less reliable data, we re-ran the analysis excluding the lowest retirement savings bracket (zero to $24,999), and noticed for these participants the across year correlations were higher at .583 (Spearman .559), and with the Ghiselli et al correction .778. An area for further research would be to investigate how the number of markers on the distribution builder, and the number of distributions sampled per participant, and experience with investing for retirement affect reliability.

[Insert Table 1 About Here]

Validations. Because the validations span two dependent variables (outcome preferences and gamble choice) and two types of estimates of $-1/\alpha$ (from retirement distributions and small-stakes gamble distributions) across two years, Figure 5 is provided to serve as a visual guide. The goal of the validations is to show that the risk aversion parameter, as estimated by DB, is a valid predictor of preferences even in the presence of numerous traditional, sensible predictors of risk taking. In addition to the psychological and industry scales described, we look at the demographic variables age, income, and gender, which have long been studied as covariates of risk preference. In general, younger people, wealthier people, and males are found to exhibit less risk aversion (for a review see Bajtelsmit and Bernasek 2001).

[Insert Figure 5 about here]

Predicting outcome preferences. We modeled the standard deviation of the histogram chosen in the outcome preference task with six predictors. (1) $-1/\alpha$, the coefficient of relative risk aversion as estimated by DB, either in year 1 or year 2, transformed to make it amenable to regression analysis 2) ***age***, in years, which in the experiment was limited from 30 to 60 3) ***income***, log-transformed 4) ***gender***, coded 1 for male 5) ***Industrial risk profile***, as described and in the

Appendix and 6) ***Psychological risk profile,*** the participant's score on the gambling and investment subscales of the Weber-Blais-Betz Domain Specific Risk Attitude Scale (Weber, Blais, and Betz 2002) as in the Appendix.

[Insert Table 2 about here]

Regressions 1- 4 in Table 2 show the results of predicting preferred outcomes.[10] The transformed risk aversion parameter α accounts for 22% to 65% of variance explained by the various models, and seems to have the strongest effect in tests that involve year 2. Remarkably, the DB estimate of risk aversion predicts outcome preferences expressed one year later and one year earlier than the time it is measured, and does so in the presence of 5 explanatory variables.

Regressions 5 and 6 provide a validation across tasks (gambles versus retirement distributions) and years. Since the gamble distributions were only presented in year 2, there are two regressions of this type instead of four. In both cases, the risk aversion parameter is significant and explains about 56% and 33% of explained variance. No predictor is significant.

Predicting gamble choices. Regressions 7 and 8 provide a validation entirely in the much-investigated domain of small stakes gambles. Here, the risk aversion parameter predicts the number of risky gambles chosen both within year 2, and from year 2 to year 1, and accounts for 23% and 21% of explained variance, respectively. It is interesting to note that when comparing regressions 1 and 4, R^2 is higher within year 2 than year 1, and similarly, Table 1 shows greater reliability within year 2, suggestive of practice or learning effects.

Demographic analysis. Correlating the DB estimate of risk aversion with age, income, and gender, provides indirect validations. If DB provides sensible estimates, we would expect to find significant correlations (though typically low in the literature) and consistent directional relationships between these variables and the risk aversion parameter. Table 3 shows these correlations for all estimates of $-1/\alpha$ discussed so far, in addition to an average of all four estimates. Looking at both Pearson and Spearman correlations, age and income are significantly related to $-1/\alpha$ in 12 out of 16 cases ($p<.1$) spanning year of measurement and type of distribution (retirement or gamble). As with Table 1, results are shown excluding the lowest retirement savings group, and slightly stronger relationships are seen (14 out of 16 significant relationships). Notably, in all 32 cases involving age and income, the correlation is in the expected direction (age increases risk aversion, income decreases it).

[10] Note that N varies from regression to regression because not all participants in a given year are able to be predicted with data from another year. For instance, only 85 participants from Experiment 1 participated in Experiment 2. 79 of these participants provide usable parameter estimates by submitting two zero-variance distributions in a year 2.

Equally notable is that in 14 out of 16 cases, there is no relationship between gender and risk aversion, and in two cases only modest relationships are found (Pearson correlation of .179, Spearman .17). However, if 14 out of 16 cases, the expected sign is observed (being female predicts greater risk aversion). Though many studies have found relationships between gender and risk aversion, few offer ideas on why this relationship might exist. For a review and some pointers to a conceptual framework, see Bajtelsmit and Bernasek (2001).

[Insert Table 3 about here]

Summary of experiments

Using the domain of retirement investing as an example, we estimated the coefficient of relative risk aversion and the loss aversion parameter using Distribution Builder. Estimates fell within ranges observed in the literature, providing a first level of validation of the tool. Two correlation metrics were calculated within sessions and between years showing significant long-term reliability, but less so among participants who had little or nothing saved towards retirement. Stronger validation tests used DB estimates of risk aversion to predict preferred investment outcomes both within and across years and in the presence of 5 additional predictors known to correlate with risk preferences. Crossing domains, a DB designed for a small-stakes gamble task provided risk aversion estimates that significantly predicted both preferred investment outcomes and choices between gambles. Correlations with age and income provided an indirect validation of the method, and correlations with gender were for the most part not observed in these domains.

Outside of parameter estimation and validation, an interesting finding of this investigation was that the CRRA model did not fit all participants' data well. Estimating a loss aversion parameter, we found that participants who chose distributions that were not well fit by CRRA were markedly more loss averse than the rest. It appears that one group of investors, with CRRA preferences, may be satisfied with the investment advice of maintaining a constant asset allocation between a risky and risk-free asset. The other group, with loss-averse preferences, may be more concerned about the likelihood that their investments could go beneath a reference level. Interestingly, in the last years, investment and insurance firms have offered some products that offer upside gain when the market goes up, and absolute downside protection when the market goes down. The latter group of consumers might well be the intended audience for these products.

Two points are relevant in this connection. First, the investors in the experiments described here could only obtain upside gain (results greater than 75%) by accepting some downside loss (results less than 75%). Hence they were limited to strategies that would provide the reference return over a range of market outcomes with lower returns in very bad markets and higher returns in very good markets. The typical protected investment product would thus appear to be designed for investors with kinked utility curves who have more money than required to obtain their reference return with certainty. The second point is that for every investor with a strategy that responds less than 1-for-1 with market moves there must be one or more with a strategy that responds more than 1-for-1. For a discussion of equilibrium in a capital market with investors who have kinked utility curves, see Sharpe (in press).

Discussion

Constructing constructive preferences. We do not believe that consumers of investment products hold ideal probability distributions of retirement income in their heads. It may be something they have never thought about before. However, each of the millions of employees who specify a fixed asset allocation for their 401K plan is indeed signing their name to a probability distribution of wealth, a distribution of which they may never learn even the mean and standard deviation. Since some preferences need to be stated in order to open a 401K plan, consumers can benefit from a tool such as the DB to help them explore the costs and benefits of downside protection and upside gain.

Ease of use. In a paper which introduced utility theory to marketing, Hauser and Urban (1977) pointed to a drawback with the choice-between-gambles method of utility assessment, "The advantages of utility theory were obtained at a substantial cost. The measurement required a personal interview of 45 minutes and the execution of the difficult lottery questions." Distribution Builder uses real-time computation to speed up utility assessment much as Toubia and colleagues (2003) have used it to reduce the number of questions required in conjoint analysis. Some of the promise of utility theory has not been realized, as it has been too difficult, too slow, or too costly to elicit information from consumers untrained in probability. Even when such data has been obtained from a series of choices made in isolation, it has often lacked consistency (Tversky and Shafir 1992). DB exploits the speed and interactive nature of graphical, non-numerical interfaces, along with expanding Internet access to obtain information from consumers, even if they are spread about the globe, and even if their time and patience for answering questions is limited. In addition, it allows researchers to explore alternative functional forms, such as loss-averse utility functions, that would be difficult to estimate with different techniques.

Outcomes, not inputs. In the world of behavioral research, much attention is focused on choices between gambles, and not on the overall distribution that would result from considering all these gambles jointly. In the world of investment advice, the same is true—not about gambles, but about choices between individual investment funds. Risk tolerance questionnaires ask how people would feel about a single investment product that could lose varying percentages of its worth overnight, but not about the variance of a portfolio of products, which could affect well-being not just overnight, but over 30 years. The investment products in a portfolio should largely be irrelevant to the consumer. What counts are outcomes: how investments combine to give an overall risk distribution. With DB we have the first efficient method for constructing and exploring complex outcome distributions.

Process tracing. While we concentrate in this paper on the preferences revealed by the final distribution, one could in principle also assess the construction process itself, using reaction times, and recordings of the movements of markers along the outcome space. This may provide insight into identifying which parts of a probability distribution are most important to a consumer, and which they care to evaluate first. Such knowledge could, for example, inform manufacturers whether they should attend to the downside risk or upside potential in a product's design.

Financial services marketing. The presumption of efficient markets seems to render marketing irrelevant to understanding financial markets. More recent developments in behavioral finance suggest that marketing may have a much larger role. Some research in marketing has started to emphasize differences in preferences for investments (Wilcox 2003) or trading style (Dhar and Zhu 2002). Other research in behavioral finance has started to characterize gender differences in trading. These results are consistent with the basic question in marketing about differences in consumer needs. We have started to make a contribution by identifying two preference segments, one well described by standard theory, the other showing significant loss aversion. An important next step would be to examine these differences more closely and determine the correlates and causes of the differences.

Risks in health care. We have illustrated the application of Distribution Builder in a financial setting. Another particularly consequential set of consumer decisions regard health. Imagine, for example, how the Distribution Builder might be applied to choices for the management of a chronic disease. Instead of eliciting preferences for the treatments, DB could assess the preference for health states or longevity, and then use extant data to suggest the treatments that would be most likely to provide that distribution. In addition, DB could be used as a communicative tool in preventative contexts to show the risks associated with various diets, exercise regimes, or preventative medications. Other applications to health psychology are proposed in Johnson, Steffel and Goldstein (2005).

Conclusion

We describe a new method for measuring preferences over multiple-outcome risk distributions. Drawing on psychological foundations concerning frequency representations of probability, simulated experience, and integrated choices, the new technique circumvents several difficulties endemic to probability and utility assessment. We demonstrate how the method can be used to estimate the coefficient of relative risk aversion and the loss aversion parameter in a consequential financial domain. The technique passes several tests of reliability and validation in the presence of other predictors, between domains, and over time. We conclude by discussing potential future applications of the new technique.

TABLE CAPTIONS

Table 1: Reliability: Test/re-test correlations of the CRRA parameter computed: within years, between years, with correction for test-retest attenuation, and excluding those with little or nothing saved towards retirement.

Table 2:

Validation: Predicting preferred outcome distributions and gamble choices based on estimates of the CRRA parameter obtained with the DB method. The parameter accounts for 22% to 65% of explained variance, even when predicting ahead or back one year in time. Income is log of income in dollars. Alpha is –1/constant of relative risk aversion Outcome preference is the standard deviation of the participant's choice in a task that presented participants with 7 lognormal distributions of terminal wealth corresponding to a 0%, 10%, 20% … 60% investment in stock and the rest in a risk-free asset. * p<.1, **=p<.05, ***=p<.01

Table 3:

Relationship between transformed risk aversion parameter and demographic covariates age, income, and gender. * p<.1, **=p<.05, ***=p<.01

Components	**Pearson Correlation**	**Spearman Correlation**	**Count**
Test-retest distributions within Year 1	.700	.702	152
Test-retest distributions within Year 2	.803	.793	148
Between Year 1 and Year 2 averages	.431	.454	75
Above with correction for attenuation	.575	--	75
Between Year 1 and Year 2 with >$25K savings	.583	.559	57
Above with correction for attenuation	.778	--	57

Table 1

Regression #	**1**			**2**			**3**			**4**		
Predicted	Year 2 Outcome Preference			Year 1 Outcome Preference			Year 2 Outcome Preference			Year 1 Outcome Preference		
Source of $-1/\alpha$	Year 2 Retirement Dists			Year 2 Retirement Dists			Year 1 Retirement Dists			Year 1 Retirement Dists		
	Estimate	SE	t-ratio	Estimate	SE	t-ratio	Estimate	SE	t-ratio	Estimate	SE	t-ratio
Intercept	-4.072	9.841	-0.41	-5.484	14.856	-0.37	-7.230	10.891	-0.66	-8.537	12.281	-0.7
$-1/\alpha$	-52.490	8.252	-6.36***	-50.613	13.818	-3.66***	-35.643	9.029	-3.95***	-19.390	10.104	-1.92*
Age	-0.069	0.112	-0.62	0.046	0.163	0.28	-0.058	0.129	-0.45	0.079	0.142	0.56
Income	0.965	1.634	0.59	1.840	2.519	0.73	1.420	1.815	0.78	0.551	2.008	0.27
Gender	0.359	0.961	0.37	3.043	1.463	2.08**	0.669	1.089	0.61	1.653	1.212	1.36
Industry Risk Scale	0.380	0.390	0.97	0.526	0.617	0.85	0.648	0.431	1.5	0.932	0.512	1.82*
Psych. Risk Scale	0.352	0.134	2.62**	-0.032	0.189	-0.17	0.336	0.161	2.09**	0.270	0.160	1.68*
R^2	.322			.224			.199			.115		
Explained variance due to $-1/\alpha$	61.68%			64.55%			46.83%			22.08%		
# of Obs	145			79			141			136		

Table 2: Continued on next page

Regression #	**5**			**6**			**7**			**8**		
Predicted	Year 2 Outcome Preference			Year 1 Outcome Preference			Year 2 Gamble Choice			Year 1 Gamble Choice		
Source of $-1/\alpha$	Year 2 Gamble Dists			Year 2 Gamble Dists			Year 2 Gamble Dists			Year 2 Gamble Dists		
	Estimate	SE	t-ratio	Estimate	SE	t-ratio	Estimate	SE	t-ratio	Estimate	SE	t-ratio
Intercept	12.028	9.703	1.24	-2.575	14.075	-0.18	1.844	0.708	2.61**	1.692	.739	2.29**
$-1/\alpha$	-21.622	6.536	-3.31***	-16.514	9.249	-1.79*	-0.977	0.477	-2.05**	-0.939	.498	-1.89*
Age	-0.130	0.128	-1.02	-0.009	0.174	-0.05	0.002	0.009	0.26	-0.025	.010	-2.6**
Log inc	1.926	1.878	1.03	3.950	2.749	1.44	-0.142	0.137	-1.04	0.104	.143	0.73
Gender	1.150	1.097	1.05	1.699	1.513	1.12	0.295	0.080	3.69***	0.093	.083	1.11
R^2	.124			.117			.120			.109		
Explained variance due to $-1/\alpha$	56.40%			33.05%			22.53%			21.32%		
# of Obs	142			78			142			142		

Table 2: Continued.

Condition	Variable	Correlate	Pearson	p	Spearman	p	N
All Participants	Y1&Y2 Alpha	Age	.286	.013 **	.261	.024 **	75
		Income	-.215	.066 *	-.109	.354	74
		Gender	-.013	.909	-.001	.991	75
	Year 2 Alpha	Age	.182	.028 **	.200	.016 **	146
		Income	-.062	.456	-.081	.335	145
		Gender	.022	.794	.040	.630	146
	Year 2 Gamble Alpha	Age	.238	.004 ***	.211	.011 **	143
		Income	-.170	.042 **	-.123	.145	142
		Gender	.071	.397	.040	.639	143
	Year 1 Alpha	Age	.175	.040 **	.154	.070 *	139
		Income	-.174	.043 **	-.159	.065 *	136
		Gender	.056	.512	.074	.389	139
>$25K Retirement Savings	Y1&Y2Alpha	Age	.301	.030 **	.261	.062 *	52
		Income	-.333	.016 **	-.269	.054 *	52
		Gender	.042	.769	.094	.506	52
	Year 2 Alpha	Age	.243	.015 **	.238	.017 **	100
		Income	-.190	.059 *	-.214	.033 **	100
		Gender	.135	.181	.147	.144	100
	Year 2 Gamble Alpha	Age	.313	.002 ***	.268	.008 ***	98
		Income	-.195	.055 *	-.186	.067 *	98
		Gender	.179	.079 *	.170	.094 *	98
	Year 1 Alpha	Age	.131	.208	.127	.222	94
		Income	-.249	.017 **	-.234	.025 **	92
		Gender	.065	.534	.076	.467	94

Table 3

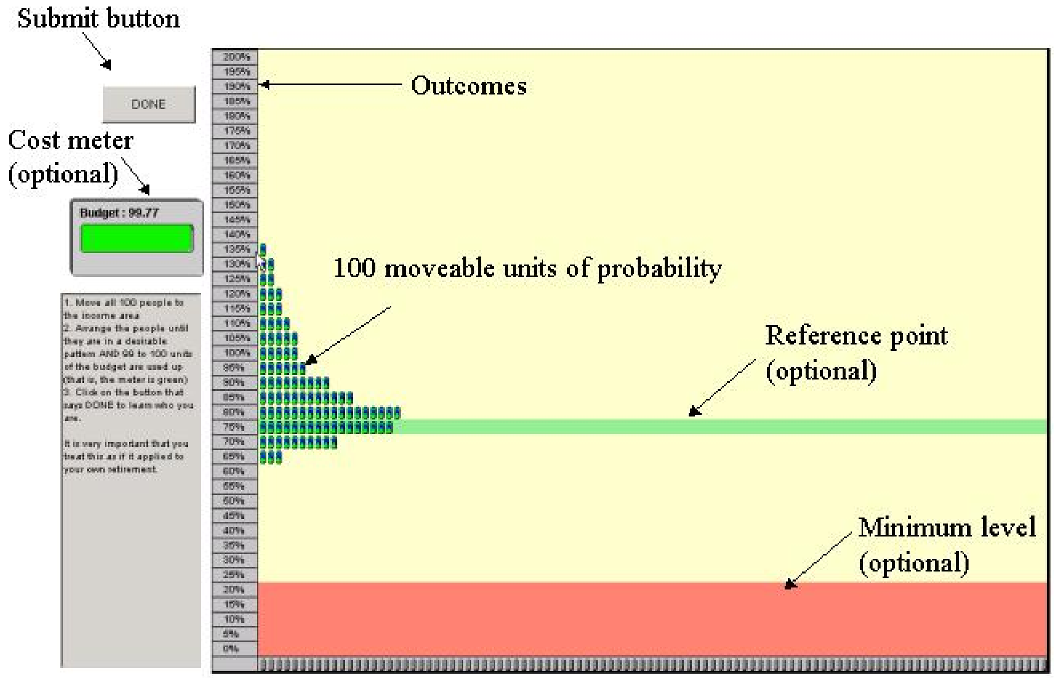

Figure 1: The Distribution Builder interface. Using moveable units of probability, participants can create arbitrarily-shaped discrete probability distributions over numerous outcomes (on the vertical axis). Between 2 and 40 outcomes and 1 and 100 units of probability can easily be displayed on a standard-size monitor. The 40 outcome / 100 unit case provides over 10^{34} unique distributions to choose among. A cost meter (upper left), can be used to constrict the space of allowable distributions, for example, to those that have a particular risk-return relationship. The cost meter functions by not allowing one to submit a distribution (using the 'submit' button on the upper left) until it satisfies an arbitrary cost function. Users can see how every change to the distribution affects the cost meter numerically and graphically. All movements are seen in the context of their effects on the system as a whole.

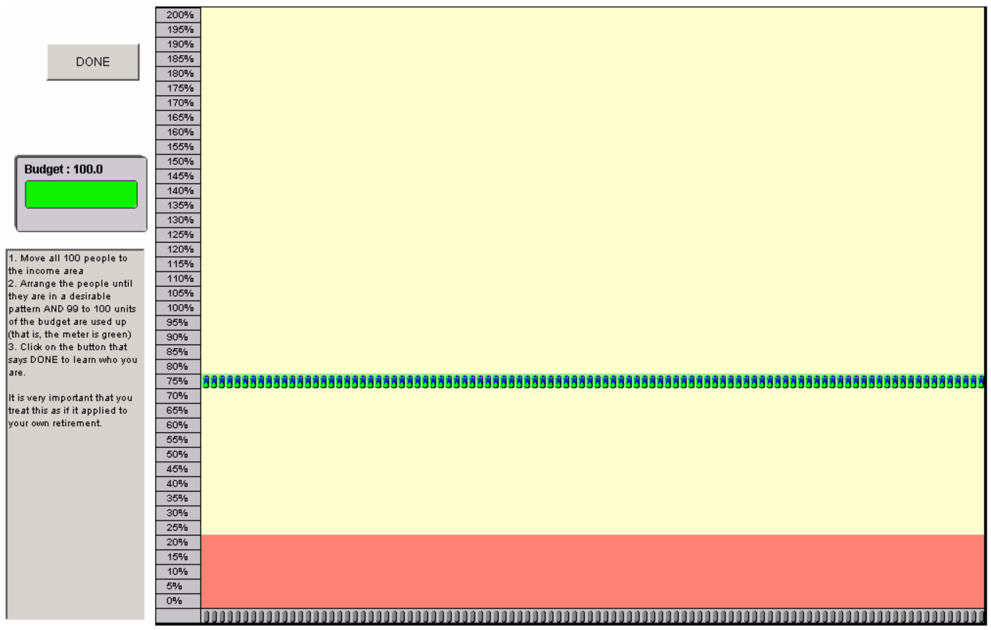

Figure 2: The least-variance distribution that satisfies the cost meter has all 100 markers on the 75% row. Notice that Figure 1 has about the same cost but a visibly higher expected value, owing to the risk-return consequences of a declining cost vector.

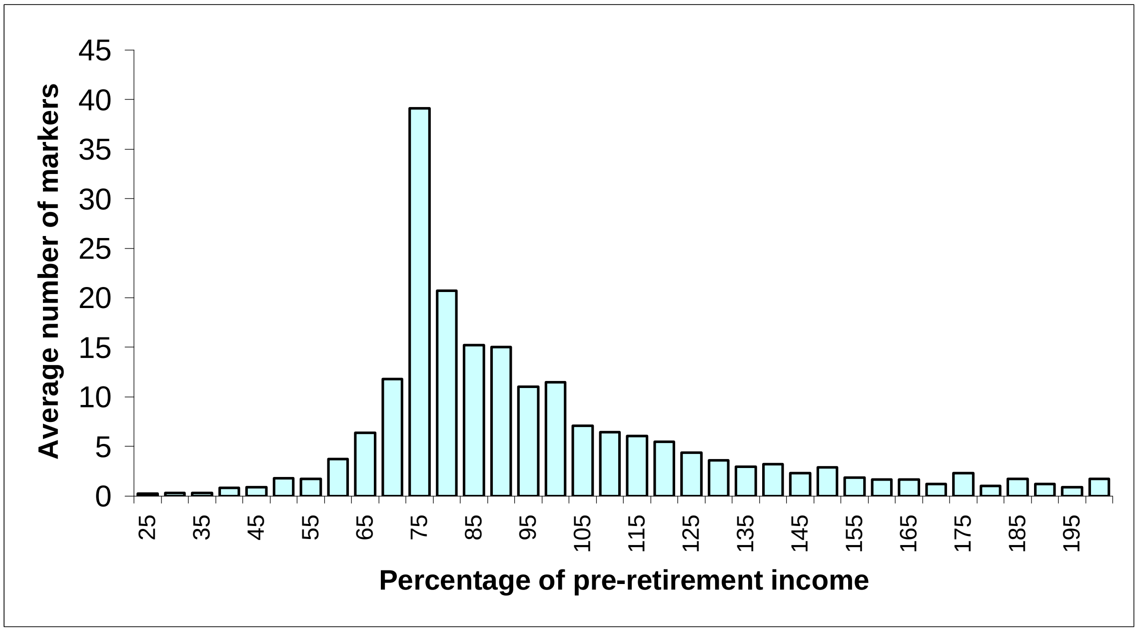

Figure 3: The distribution of the average investor, created by averaging the number of markers at each wealth level across all subjects. The distribution is right skewed and shows a peak of twice the next highest level at the reference point. 75% served as a reference point because participants were told it is a typically recommended goal level for income in retirement, and because 75% was the outcome that could be obtained without taking any risk (placing all 100 markers at 75% would satisfy the budget meter).

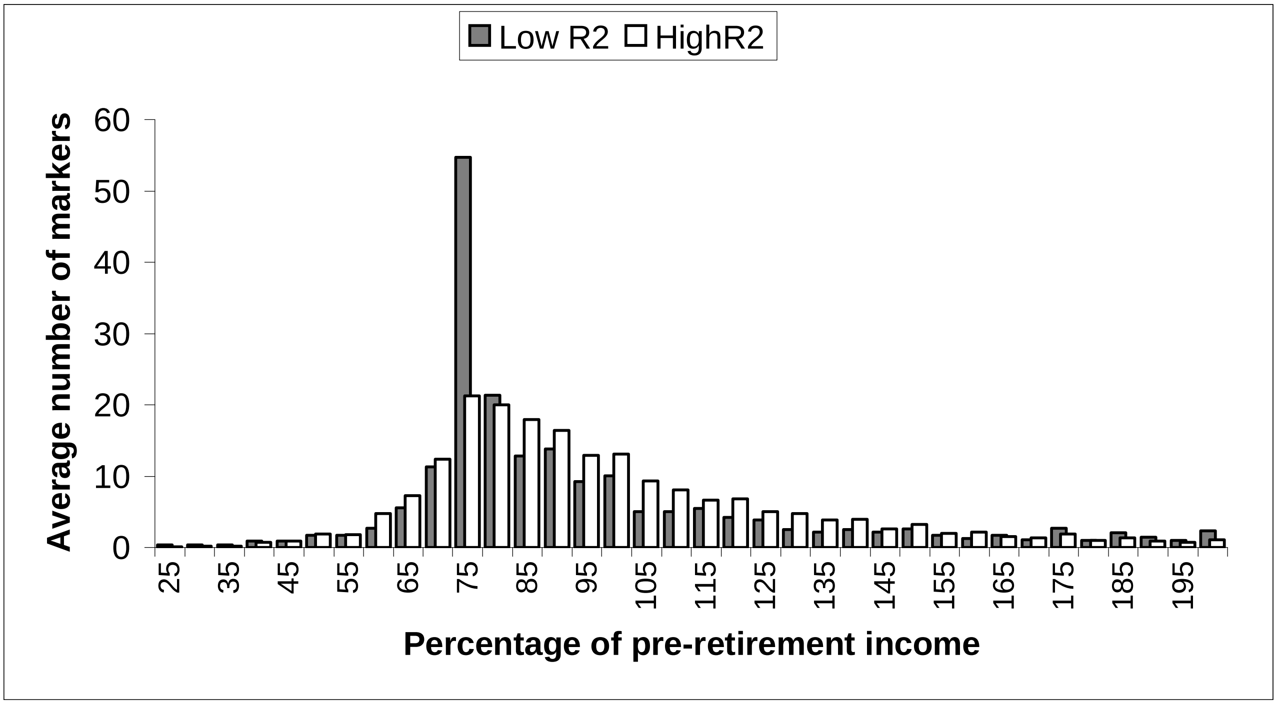

Figure 4: Composite distribution of participants who are well (high R^2) and poorly (low R^2) fit by the CRRA model. The low R^2 people do not deviate arbitrarily, but rather in a systematic way, from CRRA (lognormality). The massing of probability at the reference point (here, 75%) is a property of a loss-averse distribution.

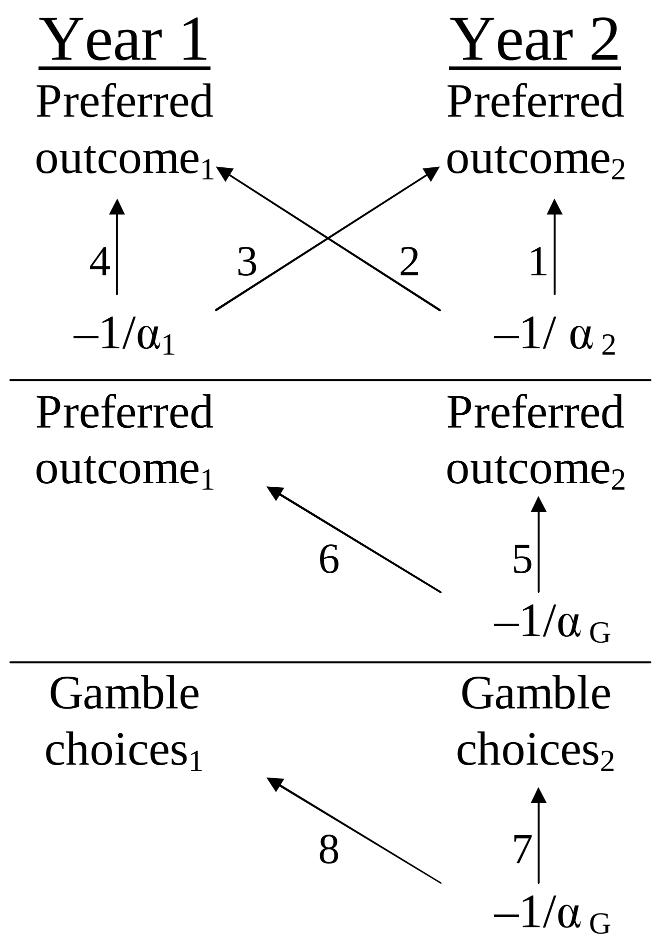

Figure 5: Guide to the 8 validation regressions.

APPENDIX

Psychological risk-taking scale used in validations: Investment and Gambling Subscales of the Weber-Blais-Betz Domain Specific Risk Taking Scale (Weber, Blais, and Betz 2002)

Instructions. For each of the following statements, please indicate the likelihood that you would engage in the described activity or behavior if you were to find yourself in that situation.

Response alternatives. Extremely Unlikely, Moderately Unlikely, Somewhat Unlikely, Not Sure, Somewhat Likely, Moderately Likely, Extremely Likely)

- Betting a day's income at the horse races.
- Investing 10% of your annual income in a moderate growth mutual fund.
- Betting a day's income at a high-stake poker game.
- Investing 5% of your annual income in a very speculative stock.
- Betting a day's income on the outcome of a sporting event (e.g., baseball, soccer, or football).
- Investing 5% of your annual income in a dependable and conservative stock.
- Investing 10% of your annual income in a new business venture.
- Gambling a week's income at a casino.

Industrial risk-taking scale used in validations

- It's more important to preserve the money you invest than to achieve significant growth from it. 1) I agree completely 2) I agree 3) I disagree 4) I disagree completely
- I would tend to choose: 1) an investment plan that minimizes the potential for loss 2) an investment plan that allows for moderate gains and losses 3) an investment plan that maximizes potential gains regardless of the potential for loss
- Assume you had a stock fund that showed steady performance for several years, but lost 15% of its value last year. The loss is in keeping with the performance of similar funds during this time. What would you do: 1) I would sell all of my fund shares 2) I would sell some of my shares, but not all 3) I would continue to hold all my shares 4) I would buy more shares of the same fund
- The real rate of return on some investments over time can be substantially reduced by inflation. How do you feel about investment risk with respect to inflation? 1) I would like to just keep pace with inflation while minimizing my potential for loss 2) I would like to try to exceed the rate of inflation by assuming a moderate potential for loss 3) I would like to try to greatly exceed the rate of inflation by assuming significant potential for loss and high volatility
- I would be most comfortable with a portfolio if, over a one-year period, it had: 1) an expected return of 6% and low chance of losing value 2) an expected return of 10% and moderate chance of losing value 3) an expected return of 14% and high chance of losing value
- Over a three-year period, I would be happiest with a portfolio that has average yearly returns that are likely to fluctuate between: 1) 0% and 10% 2) –5% and 18% 3) –10% and 25%

Note: This scale is a variant of an actual risk attitude quiz given in marketing collateral of a major mutual fund provider. Questions have been reworded to prevent recognition.

REFERENCES

Allais, M. 1953. Le comportement de l'homme rationnel devant le risqué, critique des postulats et axiomes de l'école Américaine. ***Econometrica*** 21 503-546.

Ariely, D., J. Huber, and K. Wertenbroch. 2005. When Do Losses Loom Larger Than Gains? ***Journal of Marketing Research*** 42 (2) 134-138.

Arrow, K. 1970. ***Essays in the Theory of Risk-Bearing***. Amsterdam: North-Holland.

Arrow, Kenneth J. 1964. The Role of Securities in the Optimal Allocation of Risk-bearing. ***Review of Economic Studies*** 91-96.

Bajtelsmit, V., and A. Bernasek. 2001. Risk Preferences and the Investment Decisions of Older Americans. In ***AARP Public Policy Institute Paper***.

Barberis, Nicholas. 2000. Investing for the Long Run when Returns are Predictable. ***Journal of Finance*** 55 (1) 225-264.

Baron, J. 1988. ***Thinking and deciding***. New York: Cambridge University Press.

Bauer, R. A. 1960. Consumer Behavior as Risk Taking. In ***Dynamic Marketing for a Changing World***, edited by R. S. Hancock. Chicago: American Marketing Association.

Benartzi, S., and R. Thaler. 1995. Myopic Loss Aversion and the Equity Premium Puzzle. ***The Quarterly Journal of Economics.*** 110 (1) 73-92.

———. 2001. Naive Diversification Strategies in Retirement Saving Plans. ***American Economic Review*** 91 (1) 79-98.

Benartzi, S., and R. H. Thaler. 2001. Naive diversification strategies in defined contribution saving plans. ***American Economic Review*** 91 (1) 79-98.

Bettman, J.R. 1973. Perceived risk and its components: A model and empirical test. ***Journal of Marketing Research*** 10 184-190.

Birnbaum, M. H., and W. R. McIntosh. 1996. Violations of branch independence in choices between gambles. ***Organizational Behavior and Human Decision Processes*** 67 91-110.

Birnbaum, M. H., J. N. Patton, and M. K. Lott. 1999. Evidence Against Rank-Dependent Utility Theories: Tests of Cumulative Independence, Interval Independence, Stochastic Dominance, and Transitivity. ***Organizational Behavior and Human Decision Processes*** 77 (1) 44 – 83.

Biyalogorsky, E., and E. Gerstner. 2004. Contingent Pricing to Reduce Price Risks. ***Marketing Science*** 23 (1) 146-155.

Blake, D. 1996. Efficiency, Risk Aversion and Portfolio Insurance: An Analysis of Financial Asset Portfolios Held By Investors in the United Kingdom. ***Economic Journal*** 106 (438) 1175-1192.

Brandstaetter, Eduard, Gerd Gigerenzer, and Ralph Hertwig. in press. The priority heuristic: Making choices without trade-offs. ***Psychological Review***.

Brav, A., G. M. Constantinides, and C. C. Geczy. 2002. Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence. ***Journal of Political Economy*** 110 (4) 793-824.

Camerer, C. 2005. Three Cheers--Psychological, Theoretical, Empirical--for Loss Aversion. ***Journal of Marketing Research*** 42 (2) 129-133.

Campbell, J. Y. 1996. Understanding Risk and Return. ***Journal of Political Economy*** 104 (2) 298-345.

Christensen-Szalanski, J. J. J., and J. B. Bushyhead. 1981. Physicians' use of probabilistic information in a real clinical setting. ***Journal of Experimental Psychology: Human and Performance*** 7 928-935.

Debreu, Gerard. 1959. Theory of Value. In ***The Cowles Foundation Monograph***.

Dhar, Ravi, and Ning Zhu. 2002. Up Close and Personal: An Individual Level Analysis of the Disposition Effect. In ***Yale ICF Working Paper***. New Haven.

Dybvig, Philip H. 1988. Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market,". ***The Review of Financial Studies*** 1 (1) 67-88.

Feynman, Richard. 1967. ***The character of physical law***. Cambridge, MA: MIT Press.

Friend, I., and M. E. Blume. 1975. The Demand for Risky Assets. ***American Economic Review*** 65 900-922.

Gaba, A., and A. Kalra. 1999. Risk Behavior in Response to Quotas and Contests. ***Marketing Science*** 18 (3) 417–434.

Gigerenzer, G. 1994. Why the distinction between single-event probabilities and frequencies is relevant for psychology (and vice versa). In ***Subjective probability***, edited by G. W. P. Ayton. New York: Wiley.

———. 2002. ***Calculated risks: How to know when numbers deceive you***. New York: Simon & Schuster.

Gigerenzer, G., W. Hell, and H. Blank. 1988. Presentation and content: The use of base rates as a continuous variable. ***Journal of Experimental Psychology: Human Perception and Performance*** 14 513-525.

Gigerenzer, G., R. Hertwig, E. van den Broek, B. Fasolo, and K. Katsikopoulos. 2005. A 30% chance of rain tomorrow: How does the public understand probabilistic weather forecasts? ***Risk Analysis*** 25.

Gigerenzer, Gerd, and Daniel G. Goldstein. 1996. Reasoning the Fast and Frugal Way: Models of Bounded Rationality. ***Psychological Review*** 103 (4) 650-669.

Green, P. E., A. M. Krieger, and Y. Wind. 2004. Thirty years of conjoint analysis: Reflections and prospects. In ***Marketing research and modeling: Progress and prospects***, edited by Y. Wind and P. E. Green. Boston: Kluwer.

Hakansson, Nils H. 1970. Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions. ***Econometrica*** 38 (September) 587-607.

Hasher, L., and R.T. Zacks. 1984. The automatic processing of fundamental information: The case of frequency of occurrence. ***American Psychologist*** 39 1372-1388.

Hauser, J. R., and G. L. Urban. 1979. Assessment of Attribute Importances and Consumer Utility Functions: von Neumann-Morgenstern Theory Applied to Consumer Behavior. ***Journal of Consumer Research*** 5 251-262.

Hauser, John R., and Glen L. Urban. 1977. Direct assessment of consumer utility functions : von Neumann-Morgenstern utility theory applied to marketing. In ***Sloan Working Papers***. Cambridge, MA.

Hoffrage, U., and G. Gigerenzer. 1998. Using natural frequencies to improve diagnostic inferences. ***Academic Medicine*** 73 538-540.

Hoffrage, U., S. Kurzenhäuser, and G. Gigerenzer. 2005. Understanding the results of medical tests: Why the representation of statistical information matters. In ***Heuristics and values: Relating the universals and the particulars in thinking***

about health., edited by R. Bibace, J. Laird, K. D. Noeller and J. Valsiner. Westport, CT: Greenwood Publishing Group.

Hoffrage, U., S. Lindsey, R. Hertwig, and G. Gigerenzer. 2000. Communicating statistical information. ***Science*** 290 2261-2262.

Holyoak, K. J., and B. A. Spellman. 1993. Thinking. ***Annual Review of Psychology*** 44 265-315.

Ibrekk, H., and M. G. Morgan. 1987. Graphical Communication of Uncertain Quantities to Nontechnical People. ***Risk Analysis*** 7 (4) 519-529.

Johnson, E. J., M. L. Steffel, and D. G. Goldstein. 2005. Making better decisions: From measuring to constructing preferences. ***Health Psychology*** 24 (4) S17-S22.

Kahn, B.E., and R.K. Sarin. 1988. Modeling ambiguity in decisions under uncertainty. ***Journal of Consumer Research*** 15 265-272.

Kahn, Barbara E., and Mary Frances Luce. 2003. Understanding high-stakes consumer decisions: Mammography adherence following false alarm test results. ***Marketing Science*** 22 (3) 393–410.

Kahneman, D., and D. Lovallo. 1993. Timid choices and bold forecasts: A cognitive perspective on risk taking. ***Management Science*** 39 17-31.

Kahneman, D., P. Slovic, and A. Tversky, eds. 1982. ***Judgement under uncertainty: heuristics and biases***. Cambridge: Cambridge University Press.

Kahneman, D., and A. Tversky. 1979. Prospect Theory: An Analysis of Decisions Under Risk. ***Econometrica*** 47 263-291.

Kalra, A., and M. Shi. 2001. Designing Optimal Sales Contest: A Theoretical Perspective. ***Marketing Science*** 20 (2) 170-193.

Kivetz, R. 2003. The effects of effort and intrinsic motivation on risky choice. ***Marketing Science*** 22 (4) 477–502.

Koehler, J. J. 1996. The base rate fallacy reconsidered: Descriptive, normative and methodological challenges. ***Behavioral and Brain Sciences*** 19 1-53.

Locander, W. B., and P. W. Hermann. 1979. The effect of self-confidence and anxiety on information seeking in consumer risk reduction. ***Journal of Marketing Research*** 21 268-274.

Loomes, G., and R. Sugden. 1998. Testing different stochastic specifications of risky choice. ***Economica*** 65 581-598.

Lopes, L. L. 1987. Between hope and fear: The psychology of risk. ***Advances in Experimental Social Psychology.*** (20) 255-295.

Lopes, L. L., and G. C. Oden. 1999. The role of aspiration level in risky choice: a comparison of cumulative prospect theory and SP/A theory. ***Journal of Mathematical Psychology*** 43 286-313.

Luce, R. D., and J. W. Tukey. 1964. Simultaneous conjoint measurement: A new type of fundamental measurement. ***Journal of Mathematical Psychology*** 1 1-27.

Manis, M., I. Dovenalina, N. E. Avis, and S. Cardoze. 1980. Base rates can affect individual predictions. ***Journal of Personality and Social Psychology*** 38 287-298.

Mehra, R., and E. C Prescott. 1985. The Equity Premium: A Puzzle. ***Journal of Monetary Economics*** 15 (2) 145-161.

Merton, Robert C. 1969. Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case. ***Review of Economics and Statistics*** 51 247-257.

———. 1971. Optimum Consumption and Portfolio Rules in a Continuous Time Model. ***Journal of Economic Theory*** 3 373-413.

Novemsky, N., and D. Kahneman. 2005. Boundaries of loss aversion. ***Journal of Marketing Research*** 42 (2) 119-128.

Perold, Andre F., and William F. Sharpe. 1988. Dynamic strategies for asset allocation. ***Financial Analysts Journal*** 44 (1) 16.

Pratt, J. 1964. Risk aversion in the small and the large. ***Econometrica*** 32 122-136.

Quiggin, J. 1982. A theory of anticipated utility. ***Journal of Economic Behavior and Organization*** 3 323-343.

Roselius, Ted. 1971. Consumer Ranking of Risk Reduction Methods. ***Journal of Marketing Research*** 35 56-61.

Rottenstreich, Y., and A. Tversky. 1997. Unpacking, repacking, and anchoring: Advances in Support Theory. ***Psychological Review*** 104 406-415.

Rust, R. T. , J. J. Inman, J. Jia, and A. Zahorik. 1999. What You Don't Know About Customer-Perceived Quality: The Role of Customer Expectation Distributions. ***Marketing Science*** 18 (1) 77-92.

Safra, Zvi, and Uzi Segal. 1998. Constant Risk Aversion. ***Journal of EconomicTheory*** 83 (1) 19-42.

Samuelson, Paul A. 1969. Lifetime Portfolio Selection By Dynamic Stochastic Programming. ***The Review of Economics and Statistics*** 51 (3) 239-246.

Samuelson, W., and R. Zeckhauser. 1988. Status quo bias in decision-making. ***Journal of Risk and Uncertainty*** 1 7-59.

Savage, L. J. 1954. ***The foundations of statistics***. New York: Dover.

Sedlmeier, P. 1997. BasicBayes: A tutor system for simple Bayesian inference. ***Behavior Research Methods, Instruments, & Computers*** 29 328-336.

Sharpe, W. F. 2001. Individual Risk and Return Preferences: A Preliminary Survey. In ***Working paper at www.wsharpe.com***.

———. in press. ***Asset Prices and Portfolio Choice***, ***Asset Prices and Portfolio Choice***. Princeton: Princeton University Press.

Sharpe, W. F., D. G. Goldstein, and P. W. Blythe. 2000. The Distribution Builder: A Tool for Inferring Investor Preferences. In ***Working paper at www.wsharpe.com***.

Sheth, Jagdish N., and M. Venkatesan. 1968. Risk-Reduction Processes in Repititive Consumer Behavior. ***Journal of Marketing Research*** 5 307-10.

Siebenmorgen, N., E. U. Weber, and M. Weber. 2000. Communicating Asset Risk: How the format of historic volatility information affects risk perception and investment decisions. In ***University of Mannheim Sonderforschungsbereich 504***. Mannheim.

Spellman, B. A. 1993. Implicit Learning of Base Rates. ***Psycoloquy*** 4 (61) Base Rate (4).

———. 1996. The implicit use of base rates in experiential and ecologically valid tasks. ***Behavioral and Brain Sciences*** 19 38.

Sugden, R. 2003. Reference-dependent subjective expected utility. ***Journal of Economic Theory*** 111 172-191.

Thaler, R., and E.J. Johnson. 1990. Gambling with the House Money and Trying to Break Even: The Effects of Prior Otcomes on Risky Choice. ***Management Science*** 36 (6) 643-660.

Thompson, K. M., and D. L. Bloom. 2000. Communication of risk assessment information to risk managers. ***Journal of Risk Research*** 3 (4) 333-352.

Toubia, Olivier, Duncan I. Simester, J. R. Hauser, and Ely Dahan. 2003. Fast Polyhedral Adaptive Conjoint Estimation. ***Marketing Science*** 22 (3).

Tversky, A., and D. Kahneman. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. ***Journal of Risk and Uncertainty*** 5 297-323.

Tversky, A., and E. Shafir. 1992. The disjunction effect in choice under uncertainty. ***Psychological Science*** 3 305-309.

Tversky, Amos, and Craig R. Fox. 1995. Weighing Risk and Uncertainty. ***Psychological Review*** 102 269-283.

Weber, E. U., A. -R. Blais, and N. E. Betz. 2002. A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. ***Journal of Behavioral Decision Making*** 15 (4) 263-290.

Wilcox, Ronald T. 2003. Bargain Hunting or Star Gazing: Investors Preferences for Stock Mutual Funds. ***Journal of Business*** 76 (4).

Wu, G. 1994. An Empirical Test of Ordinal Independence. ***Journal of Risk and Uncertainty*** 9 39-60.