

**Risky Curves:
From Unobservable Utility to Observable Opportunity Sets**

Daniel Friedman
University of California, Santa Cruz

Shyam Sunder
Yale University

Abstract

Most theories of risky choice postulate that a decision maker maximizes the expectation of a Bernoulli (or utility or similar) function. We tour 60 years of empirical search and conclude that no such functions have yet been found that are useful for out-of-sample prediction. Nor do we find practical applications of Bernoulli functions in major risk-based industries such as finance, insurance and gambling. We sketch an alternative approach to modeling risky choice that focuses on potentially observable opportunities rather than on unobservable Bernoulli functions.

Keywords: expected utility, risk aversion, St. Petersburg Paradox, decisions under uncertainty, option theory

JEL Codes: C91, C93, D11, D81, G11, G12, G22, L83

© Copyright Revised Draft June 6, 2011

We are grateful to Nick Barberis, Donald Brown, Shane Frederick, Uri Gneezy, Werner Gueth, William Goetzmann, Glenn Harrison, Edward Kaplan, Rajneesh Mehra, Brian Mittendorf, Steven Morris, Robin Pope, Kareen Rozen, Jozsef Sakovics, Larry Samuelson, Martin Shubik, T. N. Srinivasan, Nat Wilcox, Deniz Yavuz, and participants in CEBR Conference on Measuring Risk and Time Preferences (Copenhagen 2004), Max Planck Institute at Jena, and Yale Microeconomics and Yale School of Management Faculty workshops (2011) for conversations and exchanges on the topic. We are also grateful to Kai Du for his research assistance. We alone are responsible for the paper.

Please send comments to dan@ucsc.edu or shyam.sunder@yale.edu.

Risky Curves: Do Bernoulli Functions Predict Choice?

Daniel Friedman and Shyam Sunder

It is a veritable Proteus that changes its form every instant.

-- Antoine Laurent Lavoisier (speaking of phlogiston, quoted in McKenzie, 1960, p. 91)

1. Introduction

For several decades, economists have modeled choice under risk as expected utility maximization. Here utility is represented by a curve, the graph of an increasing function of purchasing power. There are many variants on this theme, and much has been written about the exact shape of the curve in various regions, the need for a reference point to distinguish utility gains from utility losses, the need for a probability curve relating subjective to objective probabilities, whether to represent purchasing power as income or as wealth, etc. But such concerns are peripheral to the scientific enterprise of prediction. The key question is: can we learn enough about the curves to use them to beat naïve extrapolation in predicting behavior in novel risky situations?

This paper briefly revisits the historical origins of received theory. After considering how to use the theory to generate testable predictions, it tours 60 years of empirical investigations. So far the harvest has been surprisingly slim, mainly because estimates of the curves are so Protean -- they shift erratically as the context changes, and exhibit little power to predict choice out of sample. We then consider the insights that curves may offer in thinking about risky industries, such as finance, insurance and gambling. The essay concludes by suggesting a largely neglected approach that might help predict risky choices.

2. The reincarnation of cardinal utility

Daniel Bernoulli (1738) conjectured that gamblers might use the concave function $u(x) = \ln x$ to evaluate a particular sort of risky bet. Jeremy Bentham (1789) used an increasing function to describe the greater happiness or utility enjoyed from consuming greater quantities of a divisible good, and argued that the function should be concave due to diminishing marginal utility. Later Marginalists (e.g. Marshall, 1890) noted that diminishing marginal utility implies downward

sloping demand curves. Early twentieth century economists such as Allen and Hicks (1934) and Samuelson (1938) successfully campaigned against such notions of cardinal utility, mainly on the grounds that the postulated functions lacked measurability, parsimony, and generality compared to ordinal measures of utility (Andreas, 2010).

At mid-twentieth century, just as the Ordinalist victory seemed complete, a small group of theorists including von Neumann and Morgenstern (1944), Arrow (1971), Friedman and Savage (1948) and Markowitz (1952), built a new foundation for cardinal utility. They proved that if a decision-maker's risky choices satisfy a short list of plausible consistency axioms, then there exists a particular utility function (or Bernoulli function, in the increasingly popular terminology of Mas-Collel et al., 1995) whose expectation those choices maximize.

This theoretical proposition launched a popular quest for empirically valid Bernoulli functions. Before summarizing the results of that quest, we offer some perspectives on how the abstract theory can generate predictions of actual human choices under risk.

A person's true Bernoulli function U is unobservable to outsiders, and perhaps is not even consciously accessible to that person. However, it is latent in any consistent set of choices. The function U maps possible consequences x (e.g., final wealth) into the real numbers. It is continuous and strictly increasing and, with mild additional technical assumptions, can be taken to be piecewise smooth (twice continuously differentiable except perhaps at a few kinks). Thus we can safely assume that $U' > 0$ almost everywhere.

The sign of the second derivative is a priori unrestricted, but the $U'' \leq 0$ case (i.e., a concave Bernoulli function) is central. On the one hand, a negative second derivative captures diminishing marginal utility, a hallmark of the older cardinal tradition. The logic is simply that a rational person will first purchase goods or service units that bring him greatest utility, before turning to other units that bring lower utility. On the other hand, by Jensen's inequality (e.g., Royden and Fitzpatrick, 2010), concavity implies risk aversion. Representing a risky situation by a non-trivial distribution of monetary outcomes, the expectation EU over that distribution of a concave Bernoulli function is less than its value $U(Ex)$ at the expected outcome. A person is deemed risk averse to the extent that the certainty equivalent of the risky situation (an x^* such that $U(x^*) = EU$) falls short of the expected value Ex ; the shortfall is called the risk premium. A linear Bernoulli function ($U'' = 0$) always has a zero risk premium, and thus represents risk neutrality.

A Bernoulli function is unique up to a normalization setting the zero point and the scale. Intrinsic measures of risk aversion are therefore normalized by the scale, or by U' . A leading measure is the coefficient of absolute risk aversion $a(x) = -u''(x)/u'(x)$. CARA, the one parameter family of Bernoulli functions with constant a , takes the form $u(x; a) = A - B e^{-ax}$, where $A, B > 0$ are arbitrary constants chosen for convenience. Another popular measure is relative risk aversion $r(x) = xa(x)$; the CRRA family $u(x; r)$ with constant $r(x) = r$ includes Bernoulli's original suggestion $\ln x$ as the special case with $r = 1$.

The space of all Bernoulli functions is infinite-dimensional and thus might seem empirically inaccessible. Fortunately, the Stone-Weierstrass theorem (e.g., Royden and Fitzpatrick, 2010) assures us that every Bernoulli function can be approximated arbitrarily closely within a well-chosen finite-dimensional parametric family. The empirical task therefore is twofold:

- (a) to estimate from an observed set of risky choices a parameter vector θ characterizing a function $u(x; \theta)$ that closely approximates the true Bernoulli function U , and then,
- (b) to predict subsequent behavior using the fitted function $u(x; \theta)$.

Since risky choices differ across individuals, economists soon recognized that θ might differ systematically with age, sex, nationality and other demographic characteristics. Eventually they came to recognize that θ might also vary across contexts, such as the way the risky choice was presented. We can therefore specify the estimation task (a) schematically by the equation:

$$\theta_{it} = a_0 + a_d * Demographics_i + a_i * Idiosyncrasy_i + a_c * Context_t + measurement\ error_{it} \quad (1).$$

Of course, task (a) is scientifically useful only to the extent that it improves performance in task (b) of predicting person i 's subsequent choice behavior.

3. The empirical quest, 1950-2010

The simplest empirical interpretation of the theory is that there is some particular Bernoulli function $u(x; \theta_0)$ common to all. That is, apart from minor individual idiosyncrasies and measurement error, risky choice is predicted well by a universal Bernoulli function. This function might be linear, or a member of some family such as CARA or CRRA, or perhaps

something more complicated. Markowitz (1952, Figure 4) and Friedman and Savage (1948), for example, proposed universal functions that have concave as well as convex segments.

One early study was encouraging in some respects. Edwards (1955) used a series of small and larger bets to repeatedly estimate individual utility functions for five male undergraduate students. None of the estimated Bernoulli functions departed consistently from risk neutrality, and several were consistently almost linear. Edwards used these functions, together with estimated subjective probability curves, to predict subsequent choices between pairs of bets. Due mainly to the probability curves, the predictions were far better than the naïve 50-50 prediction.

Demography. When they looked for it, later investigators typically found considerable heterogeneity across subjects' estimated Bernoulli functions. Could these differences be explained by demographics? Researchers might have hoped to make useful empirical generalizations of the form:

- Lower middle class American males of age 30 typically have Bernoulli functions well approximated in the 5-parameter Friedman-Savage family with θ near $(0, 20, 2.5, -1.2, 2.5)$, i.e., lower inflection point near income 0, upper inflection point near 20k, and CARA coefficients in the three segments of approximately $a = 2.5, -1.2, \text{ and } 2.5$.
- An upper middle income Japanese housewife of age 50 typically has a Bernoulli function approximated in the CRRA family with parameter $r = 3.0$.

Consider the field experiment reported in Binswanger (1980, 1981, and 1982). For over 100 male farmers in India, the task was to choose one of eight alternative bets of form (x_1, x_2) with $0 < x_1 \leq x_2$ and $p_1 = p_2 = 1/2$. The first alternative had no risk with $x_1 = x_2 = 50$ points, and the last alternative was $x_1 = 0, x_2 = 200$. The six intermediate alternatives were chosen so that the risk (here proportional to the payoff difference $x_2 - x_1$) has the same ordering as the expected value (here the simple average payoff $(x_2 + x_1)/2$). Binswanger repeated the task with varying stakes. His main conclusion was that the farmers tend to be more risk averse at higher stakes.

More germane to the present discussion, Binswanger (1981, Table 2) estimated the impact of demographic characteristics on his chosen risk aversion parameter, essentially $\ln(\text{CRRA})$. Wealth, schooling, age, and caste all had insignificant coefficients; only Luck had significant impact!¹

¹ Luck is defined for each subject as the number of trials where he received the higher payoff minus the number with the lower payoff. Our tentative interpretation is that farmers who win the small bets early on are more apt to choose riskier bets later.

Sillers (1980) estimated a roughly similar distribution of risk parameters for Filipino farmers. However, despite considerable effort, neither Binswanger nor Sillers found any predictive power in the fitted Bernoulli functions for the risky choice of interest—whether the farmers adopted “green revolution” techniques. Sillers (1980, p. 211) summarized his results as follows. “This chapter briefly describes an attempt to use household risk preferences, as measured in the experimental game sequence, to test the impact of household risk aversion on the rate of fertilizer applied to the dry season rice crop. This effort failed to produce a satisfactory test of the importance of this relationship or its direction... .” Studies of (male) farmers in El Salvador (Walker, 1980) and in Thailand (Grisley and Kellog, 1987) also reached negative conclusions.

Surely gender is the most prominent demographic variable that might affect Bernoulli functions. Responses to survey questionnaires consistently indicate that women on average perceive greater risk than men in a variety of personal and social activities, and there is good evidence that women are less likely than men to engage in risky activities, legal and illegal. See Eckel and Grossman (2003) for a brisk summary. Of course, rather than differences in Bernoulli functions, the survey data differences might reflect mainly informational (or response bias) differences (cf., Weber, Blais and Betz, 2002) and arrest record differences might reflect mainly different opportunities. Harrison et al.’s (2002) field experiment did not reveal any differences in estimated risk attitudes by gender or age.

In principle, laboratory choice data can isolate the impact of gender on Bernoulli functions. Many of the dozens of relevant studies seem to corroborate the conventional view that women tend to be more risk averse than men. Powell and Ansic (1997), for example, report that their female subjects had less negative risk premiums (i.e., were less risk seeking) in laboratory tasks than the male subjects. However, there are also several laboratory studies that reach different conclusions. In particular, Schubert et al. (1999) find that women subjects on average are more risk averse in abstract gambling tasks in the gain domain, less risk averse in the loss domain, and not consistently different from men in context-rich tasks in either domain. They conclude:

Our findings suggest that gender-specific risk behavior found in previous survey data may be due to differences in male and female opportunity sets rather than stereotypic risk attitudes. Our results also suggest that abstract gambling

experiments may not be adequate for the analysis of gender-specific risk attitudes toward financial decisions. [p. 385].

Table 1 in the Eckel and Grossman (2003) survey lists 24 findings from the literature, only half of which corroborate the conventional view; the others conclude that there are no systematic differences or that men are more risk averse. The authors conclude that the evidence is inconsistent, perhaps due to differences across studies in task details.

Wealth and age are often thought to be correlated with risk preferences, but there is little supporting evidence. Harbaugh et al. (2002) find young children's choices are consistent with under-weighting the low probability events and over-weighting the high probability events. This tendency diminishes with age and disappears among adults; age has no other discernable impact on risk preferences.

The literature contains some scattered results regarding ethnicity. Zinkhan et al. (1991) found that their Spanish subjects were more willing to take risks than the Americans. Harrison et al. (2003) report a field experiment in Denmark that showed no age or gender effects but indicated an education effect. Henrich and McElreath (2002) directly estimated the risk preferences of two groups of small-scale farmers (Mapuche of Chile and Sanghu of Tanzania) and, surprisingly, found them to be risk-preferring decision makers. Sex, age, land holdings, and income did not predict risk preferences and wealth was only marginally predictive. The authors note that these tribal people rarely engage in cash transactions, and conjecture that gambles in more familiar currencies such as livestock might yield different conclusions.

Yook and Everett (2003) used investment company questionnaires with MBA students to assess their risk tolerance and risk capacity scores and found that age and gender played no role in explaining their portfolio held in stocks. The income variable loaded significantly, but that may have been due to their definition of risk tolerance.²

Leland and Grafman (2003) report a surprisingly negative result. They compare normal control (NC) subjects to others who had brain damage in the ventromedial prefrontal cortex (VM). Earlier studies had found large performance differences in one

² Questionnaires from Investment Technologies, A.G. Edwards & Sons, William Droms, Scudder Kemper, Fidelity and Vanguard were used in the study. On Vanguard's website (Flagship2.vanguard.com as of May 26, 2004), risk tolerance is defined as "An investor's ability or willingness to endure declines in the prices of investments while waiting for them to increase in value." This measure seems likely to be a function of income and wealth.

complicated risky task, and the standard interpretation is that the VM brain structures are involved in making risky choices. However, these authors found no significant differences between the two groups for any of their simple risky tasks, and cite other studies that yielded mixed findings. The authors conjecture that VM brain damage affects the way people engage in a task and respond to feedback, but does not affect risk preferences per se.

Harrison and Rutström (2008) apparently is the most comprehensive survey of recent laboratory experiments covering the impact of demographics on risk aversion. They argue that the most reliable instrument for measuring risk attitudes is the Holt-Laury multiple price list (MPL),³ and report data from two large recent studies using the MPL in their Tables 2 and 3, and report another recent study in Table 4. Table 2 indicates that, among the 181 subjects tested at Georgia State University, the Hispanic subjects on average have somewhat lower coefficients of relative risk aversion, but other ethnicities (Black and Asian) have no significant impact. Nor does gender, age or marital status, or any other demographic variable, with three minor exceptions. Compared to faculty and staff, students have slightly higher CRRA, and so do subjects with upper-middle household income, compared to all other groups with higher or lower income. Compared to other majors, business majors had marginally significantly ($p = 0.05$) lower CRRA.

Table 3 considers 178 student subjects (at University of South Carolina) and finds no significant impact of gender, age, ethnicity, major, college year, grade point average, or parental education. US citizens were marginally significantly more risk averse. By far the most important impact on measured risk aversion was whether the high stakes treatment came before or after the lower stakes treatment. Table 4 considers 156 adults of various ages, mostly in Oregon. None of the demographic variables had consistent significant impact, although for one subset of tasks (gain domain lotteries), measured CRRA increased with age at first and then tapered off, while in the loss domain there was only a marginally significant ($p = 0.07$) interaction effect of age and gender.

³ In each row j , the subject chooses between lottery A and lottery B. Lottery A has two prizes close together, e.g., $x_{AH} = \$2$ and $x_{AL} = \$1.60$, while the prizes are more extreme in lottery B, e.g., $x_{BH} = \$4$ and $x_{BL} = \$0.10$. The same probability p_i of getting the larger prize applies to both lotteries in each row, e.g., $p_1 = 0.10$ in the first row, ..., $p_9 = 0.90$ in the ninth row, and $p_{10} = 1.00$ in the last row. Virtually all subjects will choose lottery A in the first row and lottery B in the last row. A subject who switches to B in the 4th row, for example, is revealed less risk averse (or more risk seeking) than one who switches in the 5th or 6th row.

To summarize, measured risk aversion is negative for some tribesmen, and it may be slightly lower for men, business majors and some Hispanic groups, but the effects are neither large nor consistent across measurement instruments or risky choice tasks. All demographic characteristics examined in the literature we have found --- including gender, age, income, wealth, and ethnicity --- have a much smaller systematic effect than one might have supposed. Indeed, most published studies (and one wonders how many studies that never saw the light of day) are unable to reject the null hypothesis that the demographics impact coefficient vector a_d in equation (1) is zero. Hopes now seem groundless that demographic generalizations (such as for the hypothetical Japanese housewife mentioned earlier) will ever provide useful predictions.

Idiosyncrasy. A possible reason for the negative results might be that individual Bernoulli functions are largely idiosyncratic, analogous to blood types. Age, gender, and wealth can't predict whether a person is A-positive or O-negative, but a single test of blood type gives an extremely accurate prediction of reactions to blood transfusions and of subsequent blood test results. By analogy, perhaps we can get predictive power at the individual level, so that risk measurements using one instrument might help predict measures with other instruments or (more importantly) behavior in new risky tasks. In terms of equation (1), the question is whether the coefficient a_i is nonzero and is useful for predicting choice behavior.

This question is different than the group-level stability question that is more often addressed in the literature. It is not without interest to find essentially the same distribution of parameter estimates for a particular group of subjects given the same task on a different day. But even when such a result holds, it merely suggests that naïve extrapolation of behavior on the first day should also predict well the behavior on the second day, without any benefit from the intermediate step of fitting a Bernoulli function on the first day's data. On the other hand, a shift in the distribution has implications regarding predictability. For example, Harrison and Rutstrom (2008, p.84-85) note a shift towards risk neutrality in the distribution of estimated relative risk aversion coefficients when the expected value was shown explicitly in an otherwise standard risky choice task. This implies that the estimated coefficient of at least some subjects changed in response to an inconsequential change in the task.

The crucial tests here measure risk preference of the same subject in several different ways. For over 100 subjects, Harlow and Brown (1990) compared four different risk attitude measures: (a) CRRA estimated from bids in first price auctions with independent private values,

(b) responses to the widely used MMPI survey questions, (c) responses to another psychometric survey, SSSV, and (d) a physiological measure (platelet monoamine oxidase or MAO concentration) known to correlate with the psychometrics. They found weak but significant correlations for male subjects between (a) and the other measures,⁴ but no relation for female subjects.

A brief excursion on risk aversion and laboratory auctions may be instructive. Observed bids in first-price sealed-bid independent-private-value auctions are typically higher than in equilibrium derived from assuming risk neutrality. Risk aversion can account for such overbidding, as worked out most carefully in CRRAM (Cox et al., 1988, the model used in (a) by Harlow and Brown). Ockenfels and Selten (2005) challenge this explanation, and show that the steady state of a plausible adaptive process (IBE, or impulse balance equilibrium”) also can explain overbidding. IBE can also explain the effect of information treatments, but risk aversion can only do so if the information treatment for some unanticipated reason were to shift risk parameters in just the right way. In theory, risk aversion leads to lower bids than does risk neutrality in third price auctions. Kagel and Levin (1993) find that actual bids indeed tend to be lower than the risk neutral benchmark when there are only 5 bidders, but tend to be above the benchmark, suggesting risk-seeking, when there are 10 bidders.

Returning to studies that track individual subjects across tasks, Isaac and James (2000) found a strong *negative* correlation between risk aversion as measured in a first price auction and risk aversion for the same individuals as measured via the traditional Becker-DeGroot-Marschak mechanism. The separate measurements corroborate earlier studies, so no additive bias correction, nor even a monotone transformation, can account for the inconsistent measurements across tasks. Berg et al. (2005) report a similar negative result, and offer the comment: “...Such a result leads to the difficult problem that there simply might not be such things as (risk) preferences...” (p. 4213).

Kachelmeier and Shehata (1992) infer risk-seeking preferences when their subjects sell a gamble, and infer risk-averse preferences when the same subjects buy the gamble. Several studies, including and Berg et al. (1992), and Fong and McCabe (1999) find that the nature of

⁴ Even this positive result is undercut by the fact that estimates of two other parameters (with no theoretical relation to risk) have correlations with (b), (c) and (d) of about the same level of significance.

(and personal involvement in) the task and institution and substantially affects measured risk aversion.

At least one field study reports more encouraging results. Harrison et al. (2004) conclude that the estimated risk attitudes in Denmark vary across identifiable populations and find significant deviation from risk neutrality for sufficiently large lotteries. Unlike Binswanger and Sillers, Harrison et al. do not report attempts to use the estimated risk attitudes to predict choices in real life out-of-sample risky situations.

Other field studies reach negative conclusions. Barseghyan et al. (2011) fit a structural model including Bernoulli functions to household data on home and car insurance decisions. They reject the hypothesis of stable risk preferences; for example, a typical household exhibits greater risk aversion in their home deductible choices than in their auto deductible choices. They conclude that “unobserved heterogeneity is not a plausible explanation” (p. 593) and “...our results call into question the empirical validity of the assumption of context-invariant risk preferences and caution against extrapolating estimates of risk preferences across contexts.” (p. 622).

To summarize, our search of the empirical literature up to 2010 suggests that demographics have very little impact on parameter estimates θ , and that individual idiosyncrasies seem unstable and often shift unpredictably (or even reverse themselves) across tasks. We do not yet seem able to identify regularities of Bernoulli functions, much less gained hope that we can predict risky choices in new tasks using estimated functions.

In fairness we should add that only recently have investigators focused on the crucial standard of out-of-sample prediction. Wilcox (2011) reports first results of a research program that could potentially yield more positive findings. He observes 100 risky choices per day on three consecutive days by 80 subjects, and uses 2/3 of the data to fit Bernoulli functions (together with subjective probability curves and a decision noise parameter). Wilcox proposes a normalization to deal with one sort of context effect, the width of the payoff range. He finds that statistical power is surprisingly low, but his design and econometric techniques are sufficiently strong to demonstrate that the normalization improves predictions of the 1/3 of the data not used in the estimations. Wilcox has not yet demonstrated that the best of his models can outpredict naïve extrapolation or can predict well in tasks that differ appreciably from the task used for parameter estimation.

4. Risk-based Industries

Even if it is not possible to predict individual acts of choice from estimated Bernoulli functions of individuals, it is still possible that these functions yield valuable insight into important macro-level phenomena such as stock and bond markets, insurance, and gambling.

Stock Market. Markowitz (1952) extended the logic of Bernoulli functions to construct a theory of how investors should select stock portfolios, and Sharpe (1964) and Lintner (1965) elaborated the equilibrium implications. In particular, they predict a positive linear relation between the expected return on any asset and its incremental risk in a diversified portfolio; the slope coefficient is called β . Unfortunately, after some initial success, the prediction has fared poorly in empirical work (see Figure 1). Leading authorities conclude:

Like Reinganum (1981) and Lakonishok and Shapiro (1986), we find that the relation between β and average return disappears during the more recent 1963-1990 period, even when β is used alone to explain average returns. The appendix shows that the simple relation between β and average return is also weak in the 50-year 1941-1990 period. In short, our tests do not support the most basic prediction of the SLB (Sharpe-Lintner-Black) model, that average stock returns are positively related to market β s. (Fama and French, 1992, p. 428).

Brealey and Myers (1996) brazenly shift the burden of proof to those who may question the theory: “What is going on here? It is hard to say. ...One thing is for sure. It will be very hard to reject the CAPM beyond all reasonable doubt.” (pp. 187-8). Whatever the source of these empirical difficulties—and many sources have been suggested—portfolio theory can no longer be counted among the success stories for the standard theory of risky choice.

Mehra and Prescott’s (1985) equity risk premium puzzle presents another serious problem for existing theory. Reasonable calibrations suggest that the stock market returns on average should carry a premium of about 0.5 percent above the returns on safest assets, but historical premiums average about 10 times this amount; see Mehra 2003, Tables 1 and 2. That paper concludes: “It underscores the failure of paradigms central to financial and economic modeling to capture the characteristic that appears to make stocks comparatively so risky.”

Bonds. The familiar bond ratings—by Standard & Poor, Moody’s, and Fitch—are a matter of judgment by experts, and reflect mainly their assessment of the chances that the borrower will default on the payment of coupons and/or the principal. They are not based on the

dispersion of outcomes that lead agents with concave Bernoulli functions to demand risk premiums. Of course, investors are ultimately concerned with bond prices and yields, and they too are largely a function of default expectations and liquidity. Even with risk-neutral investors, one expects to see a higher promised yield on lower-rated bonds simply because their holders must be compensated for accepting a higher expected default rate. Thus the higher yields to maturity on low-rated bonds cannot be taken as prima facie evidence that bondholders have concave Bernoulli functions.

Insurance. The negative actuarial (i.e., expected) value of insurance policies is often cited as evidence of widespread risk aversion. The mere existence of a vast insurance industry,⁵ the usual argument goes, demonstrates the preponderance of concave Bernoulli functions over the outcomes of insured events.

We will argue in section 6 below that there are better explanations, and show that even risk neutral people have good reason purchase standard insurance policies. For now, we simply note that, with insurance as with bonds, the relevant risk consideration is the possibility of loss, not the dispersion of outcomes. In standard theory, risk-averse decision makers dislike positive deviations from the mean as much as they dislike negative deviations. The preponderance of downside insurance, and virtual absence of upside insurance, suggests to us that the usual argument may be missing something.

Gambling. Some eighty percent of US adults report having engaged in gambling at some time in their lives, and a significant minority are heavy gamblers. The gambling industry is surely large and pervasive enough to deserve theoretical attention.⁶

Just as economists invoke concave Bernoulli functions to explain insurance, they invoke convexity to explain gambling—for a convex Bernoulli function, the certainty-equivalent is larger than the expected value, making some negative expected value gambles acceptable. Indeed, since a mean-preserving spread always increases the expectation of a strictly convex Bernoulli function, a rational person with such preferences will, at a fixed degree of actuarial unfairness, always seek the largest bet possible. A Markowitz-type Bernoulli function predicts a

⁵ In 2009, the insurance industry collected \$4,066 billion in premiums world-wide, and \$1,239 billion in the US and Canada alone, accounting for a significant fraction of the economy (Swiss Re 2010, p. 14).

⁶ The industry had gross revenue (amounts wagered less the amount paid to bettors) of \$92.27 billion worldwide in 2007 (American Gaming Association www.americangaming.org/about/overview.cfm). If revenue was about 10 percent of the bet on average, the amounts wagered would be in the neighborhood of a trillion dollars. This is

preference for gambles with an infinite downside, because the convex domain has no lower bound. Preferred bets for a Friedman-Savage Bernoulli function u , convex over only a finite interval $[b, c]$, are also much more extreme than one might think. John M. Marshall (1984) shows that the optimal fair bet (or the optimal bet with a moderate degree of unfairness) involves only two possible outcomes a and d such that $u'(a) = u'(d)$. As shown in Panel F of Figure 2, these tangency points lie beyond a and b , the two Bernoulli function's inflection points, i.e., $a < b < c < d$, so the optimal bet is quite large. Also contrary to common sense, the model predicts that over the convex domain the person will always prefer uncertainty over certainty, and at any time of day or night is willing to pay to obtain a fair (or moderately unfair) gamble.

Gambling has provoked a considerable body of research. Most studies regard the monetary consequences as important but not the only factor relevant to gambling behavior. Maximizing the expectation of a Bernoulli function, however complicated, accounts only for the monetary consequences. It ignores the thrill, the hormones, the heart rate and arousal, the bluff, the competition, and the show off (see Pope [1983], Anderson and Brown [1984], Wagenaar [1988], and McManus, [2003]). We could not find any attempts to empirically isolate the monetary and non-monetary consequences of gambling. Extended discussions of the psychology of gambling can be found in Michael B. Walker (1992), Gudgeon and Stewart (2000) and <http://www.chass.utoronto.ca/~johnbell/Final/possessionritual.html>.

Summary. Neither stock nor bond market data provide much empirical support for concave Bernoulli functions. The rapid expansion of gambling across the world doesn't either. We will soon show that insurance can largely be explained by analyzing opportunity sets under risk neutrality. Empirical support for Bernoulli functions in macro phenomena thus seems to be as scarce as in micro-level observations gathered from laboratory and field.

5. Looking Backward

The concept of phlogiston, first suggested by Greek philosophers, entered the scientific mainstream with the work of Georg Ernst Stahl (1660-1734). Postulated as an invisible compressible fluid that carried heat from one object to another, phlogiston appealed to intuition and seemed able to organize some disparate physical phenomena such as combustion of charcoal

consistent with the estimated \$550 billion wagered annually in organized gambling a decade earlier (National

(it released phlogiston) and smelting of metal ores (the metal absorbed phlogiston). However, the concept produced vexing puzzles and few novel predictions. The fluid was never isolated in the laboratory. After the emergence of Lavoisier's powerful oxidation/reduction theory in the late 1780s, phlogiston theory faded away (McKenzie, 1960, chapter 6).

Is the Bernoulli function a 20th century analogue of phlogiston? It is the centerpiece of a theory of risky choice anchored by an elegant mathematical representation theorem. Marrying Marshallian diminishing marginal utility to risk aversion enhances its appeal.⁷ Although students often find the theory unintuitive at first, it grows on them and eventually dominates their thinking as they become immersed in the discipline. There is only one problem: the theory has not yet delivered the promised empirical goods. Sixty years of intensive search by theorists and empiricists in economics, game theory, psychology, sociology, anthropology and related disciplines has not yet produced evidence that assuming people to have Bernoulli functions can help predict their risky choices. Nor does the idea seem to have helped industry practitioners.

Phlogiston theory did not disappear when it encountered puzzles, such as having a positive mass in charcoal and some metals such as magnesium, but a negative mass in other metals such as mercury. Its proponents constructed elaborate defenses, reminiscent of pre-Copernican epicycles explaining the movements of planets. Phlogiston did not vanish from respectable science until a better theory came along.⁸ Even if the lack of supporting evidence is acknowledged, expected utility theory will survive until economists are convinced that they have something better to replace it.

What might that be? Some economists regard Kahneman and Tversky's (1979) prospect theory as a leading candidate. We do not share that view. Prospect theory strikes us as an especially flexible variation on the standard theme. It postulates an S-shaped value function u similar to a Markowitz (1952, Figure 4) Bernoulli function: u is convex below an inflection point

Research Council, 1999).

⁷ Since the presence of a convex segment in their proposed Bernoulli function is inconsistent with diminishing marginal utility, Friedman and Savage took pains to deny connections between the old and newer notions of cardinal utility by asserting that Bernoulli functions are "not derivable from riskless choices." (e.g., 1952, p 464). Their view does not seem to have taken hold in economics, where DMU and convexity somehow continue to coexist.

⁸ And its loyal supporters died. McKenzie remarks, "Priestley and Cavendish, on whose work much of the new theory was based, clung to the phlogiston theory to the end of their lives."

z (the reference point from which gains and losses are distinguished) and concave above.⁹ By itself, the value function predicts that people are risk-seeking in the loss domain, e.g., would not purchase insurance even at moderately subsidized prices. To explain unsubsidized insurance purchase and other inconvenient behavior, prospect theory supplements the Bernoulli function u with a probability curve w similar to that postulated in Edwards (1955) and earlier work. This flexibility (together with an unmodeled phase of editing and adjustment) allows prospect theory to rationalize a wide range of risky choice data, but we have seen no evidence that it can predict individual behavior in new risky tasks; see, among many other papers, Hey and Orme (1994) and Harless and Camerer (1994).¹⁰

6. A Way Forward

An investigator encountering difficulties should instinctively return to first principles. In that spirit we ask: What is risk?

Since 1950s, economists have equated risk with dispersion of outcomes, typically measured in terms of the second moment of the distribution. But that is not the original meaning. To older generations of economists, and to virtually all non-economists, risk refers to the possibility of harm. Dispersion matters only on the downside; the upside is not considered risky except by modern economists.¹¹ Perhaps it is time to rethink how to quantify risk.

Risk is multifaceted. Even technically sophisticated bankers distinguish operational risk from political risk and do not lump them together with counterparty risk, credit risk or market risk. The reason is that different levels and kinds of risk change the opportunity sets available to a decision maker in different ways. Our suggestion of the way forward, therefore, is to focus on how risk affects opportunity sets, rather than on how preferences interact with dispersion.

One reason for our suggestion is methodological. Traditionally, economists have distinguished themselves among social scientists by setting an austere standard for their work:

⁹ The value function has at least 3 free parameters even after specifying the reference point z and allowing for a kink there. One can normalize the right derivative $u'(z+) = 1$, but then must specify the left derivative $u'(z-) > 1$ and at least two curvature parameters, e.g., $a(x) = a_1 > 0$ for $x > z$ ["risk aversion for gains"] and $a(x) = a_2 < 0$ for $x < z$ ["risk seeking for losses"].

¹⁰ As new proposals and variations of prospect theory appear (e.g., Koszegi and Rabin 2007; Barberis and Huang 2008), we must await accumulation of empirical evidence on their ability to predict individual choice in laboratory as well as field across a range of contexts with generality and economy comparable to competing theories.

put the explanatory burden on potentially observable opportunities such as prices and incomes rather than on unobservables such as preferences or beliefs (e.g., Stigler and Becker, 1977). This maxim has often led to distinctive predictions and new insights (e.g., Stigler, 1984). Our point is that opportunity sets are potentially observable, while Bernoulli functions (and subjective probability curves) are not.

Another reason is the recent development of useful techniques. Options theory deals with one-sided phenomena and has enjoyed increasing success among financial market practitioners as well as academic researchers. As explained below, it seems useful for our purposes because some of the ways that risky choice interacts with established commitments can be described in terms of embedded real options.

We do not have a full-fledged theory to present, but instead will offer a series of illustrative examples. Consider once more the purchase of homeowners' insurance. What additional costs do people incur when they suffer losses from fire, theft, or accidents? It's not just the cost of replacement that matters, but also the time cost and aggravation of making temporary arrangements, and the increased difficulty of meeting contractual obligations. Such considerations can be captured in contingent opportunity sets, and they lead to new predictions, e.g., that homeowners with larger mortgages will carry more life insurance and less discretionary fire insurance. More generally, insurance simplifies one's life by reducing the number, diversity, and cost of contingency plans, and indirectly expands the opportunity set. It is hard to see how Bernoulli functions can capture these important considerations, or even explain life insurance.

Consider gambling. Pioneers such as Friedman and Savage (1948) thought it could be explained by convex segments of unobservable Bernoulli functions, but six decades of empirical search has not brought to light stable preferences of that sort. Indeed, Henrich and McElreath (2002) found subsistence farmers to be risk-loving in gains, while recent surveys such as Harrison and Rutstrom, 2008, p. 90ff, cast doubt that convexity exists in the loss domain. The opportunity set approach to risk would re-direct attention to potentially observable considerations such as bailout options. For example, one might predict that a low income member of a wealthy family is more likely to be a high roller because winning big would give

¹¹ Ironically Markowitz, whose portfolio theory made the dispersion measure of risk so commonplace in economics, is an exception. Markowitz (1959, p. x) proposed negative semi-variance as a measure of risk and suggested that it might provide a better approximation of an individual's utility function, albeit less convenient than variance).

him clout as well as wealth, while loosing big would only reinforce his current status without seriously threatening his survival.

For routine implementation of the opportunity set approach, we suggest distinguishing *gross payoff*, the stated values of x in the risky gamble, from *net payoff*, the ultimate change y in purchasing power arising from the gamble. For the sake of parsimony, we suggest treating the decision maker as if she were risk neutral in net payoff, and assume that she chooses so as to maximize the expected value Ey of net payoff. The linear approximation should be quite good when the DM is dealing with small-to-moderate stakes and has access to reasonably efficient financial markets. For example, a gain or loss of \$1,000 today implies a lifetime gain or loss of only a few nickels in daily consumption.¹²

It turns out that this rather simple framework can capture a wide variety of risky situations. To begin, suppose that a decision maker (DM) is endowed with some obligation $z > 0$. If he fails to meet the obligation, he faces additional costs that can be approximated as a fraction $a \in (0, 1)$ of the shortfall. For example, if the DM has a credit card balance of $z = \$1,000$ on the monthly statement and pays only \$600 by the due date, he will incur an additional cost of $\$400a$, where $a \approx .02$ is the monthly interest rate. Other obvious examples of z for household DMs include mortgage, rent, and car payments. Examples for business firm DMs include payroll obligations, debt service, and bond indentures. A biological example is the number of calories z needs to maintain normal activity; rebuilding depleted fat stores or muscle tissue incurs additional metabolic overhead of at least $a = 0.25$ and often considerably more (Schmidt-Nielsen, 1997).

Panel A of Figure 2 shows the resulting net payoff $y(x) = x - z$ for $x > z$ and $y(x) = (1+a)(x - z)$ for $x < z$. The function is concave and piecewise linear. If z is not precisely known at the time the DM makes a risky choice, e.g., if some random cash flow might partly offset the contractual obligation, then the expected net payoff $y(x)$ is strictly concave over the support of z .

Fiduciary responsibilities also lead to concave net payoffs for the trustee. When she obtains a gross payoff for the client far above the expectations, her net payoff is only slightly higher than when meeting expectations, but when the gross payoff falls short of expectations her net payoff is far lower after taking into account the legal and reputation costs. Progressive

¹² Another advantage is that linearity of the value function makes it much more plausible that the risky choice in question can be separated from the whole set of lifetime choices.

income taxes induce a similar relationship between gross and net cash flows: the slope of the function y is less at higher x due to higher marginal tax rates.

Discrete, irreversible decisions are yet another reason for concave net payoff functions. For example, suppose we see someone turn down a job offer whose expected present value clearly exceeds that of current salary plus all adjustment costs associated with the move. The usual interpretation is that this DM deducts a risk premium. Another possibility is that favorable new job offers are more likely for an established incumbent than for a new hire in a new city. Thus accepting the new job might extinguish a valuable wait option, whose value a rational DM would deduct from the new job offer. Dixit and Pindyck (1994), for example, show that deducting the value of such options leads to net payoff that is concave in the gross job offer x .

In all these cases, an uninformed outsider—one who observes only gross payoffs—will not be able to distinguish a risk-neutral DM with concave net payoffs from a risk-averse DM with a linear net payoff function. An observer with better information on net payoffs can make the distinction, and avoid the specification error of attributing an unstable concave Bernoulli function to a risk-neutral DM with varying net payoff functions.

There are also plausible circumstances that lead to specification error in the opposite direction: a risk-neutral DM can appear to be risk-seeking because his net payoff function is convex in gross payoff. A simple example is a tournament whose the only prize P goes to the DM with highest x . Assume that each of $K > 1$ contestants draws his gross payoff independently from the cumulative distribution G (obtained, for instance, in a Nash equilibrium of effort choices). Then the expected net payoff is $y(x) = PG^{K-1}(x)$, which tends to be more convex the larger the number of contestants. Panel B of Figure 2 illustrates the example for three contestants and uniform distribution G .

Business examples include decisions made in the shadow of bankruptcy, or bailout. Suppose that failure to meet a contractual obligation $z > 0$ results in bankruptcy proceedings and shortfalls are passed to creditors, as in Figure 2C. The net payoff again is $y(x) = x - z$ for $x > z$ but now is $y(x) = (1-a)(x - z)$ for $x < z$, where $a \in (0, 1)$ is the share of shortfall borne by other parties. This yields a piecewise linear convex relationship. Again presence of a random component to cash flows would smooth out the graph and make y a strictly convex function over the support of the uncertainty.

Bailouts create convex net payoffs in a similar manner. The U.S. savings and loan industry in the 1980s is a classic example. While deposit insurance was still in effect (i.e., $a > 0$), rapid deregulation made a whole new set of gambles available to these banks. The convex net payoff created an incentive to accept risky gambles in g . Indeed, some of the gambles with negative expected gross value have positive expected net value after considering the proceeds from deposit insurance; again see Figure 2C.

Certain opportunity sets would lead an uninformed outside observer to infer that a risk-neutral DM has a non-linear Bernoulli function with concave and convex segments, as suggested by Friedman and Savage (1948) or Markowitz (1952). Indeed, their intuitive justifications for these segments can be naturally re-interpreted as arising from opportunity sets that induce rational risk neutral people to behave as if they have complicated Bernoulli functions. For example, suppose the DM lives in subsidized housing with subsidy rate $a > 0$ if her income is less than or equal to z_1 , and becomes ineligible for subsidy if actual income (taking into account opportunities to disguise it) exceeds $z_2 > z_1$. If ineligible, she spends fraction $c > 0$ of incremental income on housing. Then net income (after housing) is $y(x) = y_0 + (1-a)(x-z_1)$ for $x < z_1$, and $y(x) = y_0 + (x-z_1)$ for $z_1 \leq x \leq z_2$, and $y(x) = y_0 + z_2 - z_1 + (1-c)(x-z_2)$ for $x > z_2$; see Figure 2D. After taking into account uncertainties of cash flows (or uncertainties of being caught and evicted for excess income) she would appear to have a smooth Markowitz-type Bernoulli function over gross income. Employing a perspective quite close to our own, James and Isaac (2001) derive a very similar Bernoulli function in gross payoff given a progressive tax and a bankruptcy threshold.

Friedman and Savage (1948) motivate their example with a story about the possibility of the DM moving up a rung on the social ladder. To sharpen their story a bit, suppose that z_1 is the threshold income at which the DM moves from the current working class neighborhood to a middle class neighborhood with better schools. Suppose that at a lower income z_0 the DM puts a fraction $c > 0$ of incremental income into private schools or other special expenditures that would be redundant in the new neighborhood. Finally, suppose that only at a higher income $z_2 > z_1$ does the family blend in well in the new neighborhood; at intermediate levels one has to spend a fraction $d > 0$ of incremental income on upgrading clothes, car, etc. Then, one infers a piecewise linear Bernoulli function shown in Figure 2E, which after the usual smoothing, becomes a Friedman-Savage function (see Figure 2F) with one inflection point in the interval (z_0, z_1) and a

second in (z_1, z_2) . But the characteristic non-linear shape is the result of the DM's net payoffs $y(x)$ within the available opportunity set, not some sort of intrinsic preferences.

Marshall (1984) obtains a similar shape for the indirect utility function for income. True preferences are assumed to be concave in income and increasing in an indivisible $\{0, 1\}$ good such as residential choice. He mentions other possible indivisibilities including fertility, life, and career choice. Hakansson (1970) derives a Friedman-Savage type function in an additively-separable multi-period setting. The net payoff is expected utility of wealth, given by a Bellman equation for the consumption-investment plan, assuming that the Bernoulli function of consumption each period is CRRA. The gross payoff is the present value of endowed income. He derives the desired Bernoulli function explicitly from particular constraints on investment and borrowing.

Masson (1972) drops the parametric assumptions and presents a streamlined, graphical argument in a two-period setting. Suppose the DM has standard general two-period preferences that are homothetic (hence consistent with global risk neutrality), and that consumptions at the two dates have decreasing marginal rates of substitution. Masson shows that realistic capital market constraints can create concave or mixed functions $y(x)$, where x is first period endowment and y is maximized utility in a riskless world. For example, suppose the borrowing rate b exceeds the lending rate l . Then the DM will borrow so $y' = b$ when realized x is sufficiently small, and will lend so $y' = l$ when realized x is sufficiently large. For intermediate values of x the DM consumes all of the incremental first period endowment, and $y' = \text{MRS}$, which decreases smoothly from b to l . Thus risk-neutrality in y induces Bernoulli function in x that is concave, and strictly concave over stakes such that it is not worthwhile to adjust one's bank account. Masson obtains Markowitz and Friedman-Savage type induced Bernoulli functions when the borrowing and lending rates are not constant.

Chetty (2002) derives an even more complex shape for an indirect utility function of wealth. He assumes overall concave preferences with frictional costs of deviating from a commitment to the current level of consumption decisions for one good (e.g., housing) and no such costs for the other good. The resulting net payoff function inherits from the overall function its concavity over the upper and lower extremes of gross payoff, features increased local curvature for small changes in g from the base level, joined by kinks (locally convex portions).

Finally, it may be worth revising the St. Petersburg Paradox that first inspired Bernoulli. A gamble that pays 2^n rubles with probability 2^{-n} for every $n = 1, 2, \dots, \infty$, has expected value $1+1+1+\dots = \infty$, but nobody will pay an infinite amount to play such a gamble. Bernoulli (1738) proposed that a person's willingness to play (ignoring base wealth w_0) is:

$$E \ln(.) = \sum_{n=1}^{\infty} 2^{-n} \ln(2^n) = \sum_{n=1}^{\infty} n 2^{-n} = (2^{-1}) / (1 - 2^{-1})^2 = 2.$$

We acknowledge that utility may eventually diminish, but to us a more satisfactory resolution is to note that the opportunity set of the DM is bounded. The person offering the gamble must have finite ability (and willingness) to honor a promise; above some value $2^n = B$, say, he is likely to default. Thus even a risk-neutral gambler should be willing to play no more than the expected value of the first $n(B) = \lceil \ln B / \ln 2 \rceil$ terms. In presence of upper bound $B =$ one million rubles, the willingness to pay is less than 20 rubles.

To summarize, our suggestion for a way forward to a better understanding of risky choice is careful analysis of observable opportunity sets of DMs. In particular one should identify the relationship $y(x)$ between gross and net payoffs, and see how far the simple risk neutral model can take us.

7. Concluding Remarks.

Extant theories of risky choice center on non-linear Bernoulli functions, but sixty years of empirical work has not yet made them operational. Instead of abandoning the approach, economists have proposed ever more complicated variants: utility functions with kinks, transformations via the distribution function (or rank-dependence), and subjective probability curves.

It is conceivable that such persistence will eventually pay off. Perhaps advances in econometric technique and larger scale experiments in the lab and field will isolate regularities in estimated Bernoulli functions that actually are useful in out-of-sample prediction. In terms of equation (1), this means finding stable context coefficients or idiosyncratic coefficients, or perhaps fairly simple interactions. (Of course, arbitrary interactions will not help predict out-of-sample.) Another possibility is that advances in neuroeconomics will make Bernoulli functions observable (for example, see Knutson and Bossaerts, 2007). If so, the free parameter issue would be resolved and that would become a promising way forward. So far, however, the research

suggests no simple mapping from brain processes to psychological and economic constructs, and we do not believe that a breakthrough is imminent.

Other approaches are available to applied theorists and empirical researchers who do not expect to succeed in the next few years where the previous 60 years have failed. Our own suggestion is to return to the roots of choice theory, and put the explanatory burden on potentially observable opportunities rather than on unobservable utilities and beliefs. Stigler and Becker (1977) proposed this as a general standard for economics research, and Friedman and Savage (1948) and Markowitz (1952) can be re-interpreted as nice examples of this approach. The foundations of finance are being reconstructed using options theory instead of risk aversion, and we believe parallel efforts hold great promise for other aspects of economics, including models of risky choice.

The academic literature on risky choice is vast, spread across many disciplines, and well beyond our capacity to read or review. We have tried our best (within personal constraints) to track down relevant studies and evidence, but it would be a miracle if we have not missed important and relevant pieces of work (and in the process created our own sampling bias as a matched twin of the familiar publication bias towards studies that find positive results.) We hope that you, our readers, will bring to our attention what we have missed.

References

- Allen, R. G. D. and J. R. Hicks. 1934. A Reconsideration of the Theory of Value, II," Economica, 1, n.s. 196-219.
- Anderson, G. and R. I. Brown. 1984. "Real and Laboratory Gambling, Sensation-Seeking and Arousal," The British Journal of Psychology 75:3 (August), 401-410.
- Andreas, Jonathan T. 2010. "The Cardinalist Manifesto: The Epistemology of the Measurability of Utility," Ph.D. Dissertation in the Graduate College of the University of Illinois at Chicago.
- Arrow, Kenneth J. 1971. Essays in the Theory of Risk Bearing. Amsterdam: North Holland.
- Barberis, Nicholas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices, American Economic Review 98:5, 2066-2100.

- Bentham, Jeremy. 1789. *An Introduction to the Principles of Morals and Legislation*. London: T. Payne and Son. (<http://www.econlib.org/library/Bentham/bnthPML.html>).
- Berg, Joyce, John Dickhaut, and Kevin McCabe. 2005 (1992). "Risk preference instability across institutions: A dilemma," *Proceedings of the National Academy of Sciences of the U.S.A.* Vol. 102 No. 11 March 15): 4209-4214.
- Bernoulli, Daniel. 1768. "Specimen Theoriae Novae de Mensura Sortis," Commentarii Academiae Scientiarum Imperialis Petropolitanae Tomus V [Papers of the Imperial Academy of Sciences in St. Petersburg Vol. V], pp. 175-192; English translation by Louise Sommer published as "Exposition of a New Theory of Measurement of Risk," Econometrica 22:1 (January 1954), 23-36.
- Berseghyan, Levon, Jeffrey Prince, and Joshua C. Teitelbaum. 2011. "Are Risk Preferences Stable across Contexts? Evidence from Insurance Data," The American Economic Review Vol. 101 (April): 591-631.
- Binswanger, Hans P. 1980. "Attitudes toward Risk: Experimental Measurement in Rural India," American Journal of Agricultural Economics (August), 395-407.
- Binswanger, Hans P. 1981. "Attitudes Toward Risk: Theoretical Implications of an Experiment in Rural India," The Economic Journal 91:364 (December), 867-890.
- Binswanger, Hans P. 1982. "Empirical Estimation and Use of Risk Preferences: Discussion," American Journal of Economics Vol. 64 No. 2 (May): 391-393.
- Brealey, Richard, Stewart Myers, and Franklin Allen. Principles of Corporate Finance. 4th ed. Boston: McGraw-Hill/Irwin, 2006.
- Black, Fischer. 1993. "Beta and Return: Announcement of the death of beta seem premature," Journal of Portfolio Management, pp. 8-18.
- Cox, James C., Vernon L. Smith, and James M. Walker. 1988. "Theory and Individual Behavior of First Price Auctions." *Journal of Risk and Uncertainty*, 1, 61- 99.
- Chetty, Raj. "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," Harvard University Mimeograph (2002). (<http://elsa.berkeley.edu/~chetty/papers/commitments.pdf>).
- Dixit, Avinash K. and Robert S. Pindyck. 1994. Investment under Uncertainty. Princeton University Press.

- Eckel, Catherine C. and Phillip J. Grossman, "Sex and Risk: Experimental Evidence," forthcoming in C. R. Plott and Vernon L. Smith, eds. Handbook of Results in Experimental Economics. Amsterdam: North-Holland/Elsevier Press, 2003.
- Edwards, Ward. 1953. "Probability Preferences in Gambling," American Journal of Psychology 66, 349-364.
- Fong, Christina, and Kevin McCabe. 1999. "Are decisions under risk malleable?" Proceedings of the National Academy of Sciences of the U.S.A. Vol. 96 No. 19 (September): 10927-32.
- Friedman, Milton, and Leonard J. Savage. 1948. "The Utility Analysis of Choices Involving Risk," Journal of Political Economy 56 (August), 279-304.
- Grisley, William. 1980. *Effect of risk and aversion on form decision-making*. Champaign: University of Illinois Press.
- Gudgeon, Chris, and Barbara Stewart. 2000. Luck of the Draw. Arsenal Pulp Press.
- Hakansson, Nils. 1970. "Friedman-Savage Utility Functions Consistent with Risk Aversion," Quarterly Journal of Economics 84 (August), 472-487.
- Harbaugh, William T.; Krause, Kate, and Vesterlund, Lise, "Risk Attitudes of Children and Adults: Choices Over Small and Large Probability Gains and Losses," Experimental Economics, 5, 2002, 53-84.
- Harless, David W. and Colin F. Camerer. 1994. "The Predictive Utility of Generalized Expected Utility Theories," Econometrica, Vol. 62, No. 6 (November): 1251-89.
- Harlow, W. and K. Brown. "Understanding and Assessing Financial Risk Tolerance: A Biological Perspective." Financial Analysts Journal. November/ December 1990: 50-80.
- Harrison, Glenn W. 2002. "Estimating Individual Discount Rates in Denmark: A Field Experiment," The American Economic Review, Vol. 92 No. 5 (December), 1606-1617.
- Harrison, Glenn W., Morten Igel Lau, E. Elisabet Rutström. 2004. "Estimating Risk Attitudes in Denmark: A Field Experiment," Center for Economic and Business Research and University of Central Florida Working Paper, March 2004.
- Harrison, Glenn W., Morten Igel Lau, E. Elisabet Rutström and Melonie B. Williams, "Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues," University of Central Florida Working Paper, August 2003.
- Harrison, Glenn W., and Elisabeth Rutstrom. 2008. "Risk Aversion in the Laboratory," in Risk Aversion in Experiments, Vol. 12, 41-196. Emerald Group Publishing Ltd.

- Henrich, Joseph and Richard McElreath, 2002. "Are Peasants Risk-Averse Decision Makers?" Current Anthropology 43:1 (February), 172-181.
- Hey, John D., and Chris Orme. 1994. "Investigating Generalizations of Expected Utility Theory Using Experimental Data," Econometrica, Vol 62, No. 6 (November): 1291-1326.
- Isaac, R. Mark, and Duncan James. 2000. "Just Who Are You Calling Risk Averse," Journal of Risk and Uncertainty , Vol. 20 No. 2, 177-187.
- James, Duncan, and Mark Isaac. 2001. "A Theory of Risky Choice," Working Paper (February).
- Kachelmeier, Steven J. and Mohamed Shehata. 1992. "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People's Republic of China," The American Economic Review Vol. 82 No. 5 (December): 1120-1141.
- Kagel, John, and Dan Levin. 1993. "Independent Private Value Auctions : Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders," The Economic Journal Vol. 103, No. 419 (July) : 868-879.
- Kahneman, Daniel, and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk," Econometrica, 47, 263-91.
- Knutson, Brian, and Peter Bossaerts. 2007. "Neural Antecedents of Financial Decisions," The Journal of Neuroscience 27 :31 (August 1) : 8174-8177.
- Koszegi, Botond, and Matthew Rabin. 2007. « Reference-Dependent Risk Attitudes, » American Economic Review 97 :4 (September), 1047-1073.
- Lakonishok, J. and A. C. Shapiro.1986. Systematic Risk, Total Risk, and Size as Determinants of Stock Market Returns," Journal of Banking & Finance 10, 115-132.
- Leland, J.W. and J. Grafman. 2003. "Experimental Tests of the Somatic Marker Hypothesis," National Institute of Neurological Disorders and Stroke Working Paper.
- Lintner, John. 1965. "Valuation of Risk Assets in the Selection of Risky Investments in Stock Portfolios and Capital Budgets," Review of Economics and Statistics 47 (February), 13-37.
- Markowitz, Harry M. 1952. "Portfolio Selection," Journal of Finance 7 (March), 77-91.
- Marshall, Alfred. 1890. Principles of Economics. London: Macmillan and Co. Ltd.
(<http://www.econlib.org/library/Marshall/marP.html>).
- Marshall, John M. 1984. "Gambles and the Shadow Price of Death," American Economic Review 74:1, 73-86.

- Mas-Collel, Andreu, Michael Dennis Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press.
- McManus, Jim. 2003. Murderers, Cheetahs, and Binion's World Series of Poker. New York: Farrar, Straus & Giroux..
- Masson, Robert T. 1972. "The Creation of Risk Aversion by Imperfect Capital Markets," The American Economic Review 62:1-2, 77-86.
- McCabe, Kevin and Christina Fong. 1999. "Are Decisions Under Risk Malleable?" K McCabe and Christina Fong, in the Proceedings of the National Academy of Sciences, (96), pp. 10927-10932.
- McKenzie, A. E. E. 1960. The Major Achievements of Science, Vol. I. Cambridge University Press.
- Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15 (1985), 145-161.
- National Research Council. 1999. *Pathological Gambling: A Critical Review*. Washington, D.C.: National Academies Press. <http://books.nap.edu/books/0309065712/html/R1.html>.
- von Neumann, John, and Oskar Morgenstern. 1947. The Theory of Games and Economic Behavior. Princeton, NJ: Princeton University Press.
- Ockenfels, A. and R. Selten. 2005. "Impulse Balance Equilibrium and Feedback in First Price Auctions," Games and Economic Behavior Vo. 51 No. 1, pp. 155-170.
- Pope, Robin. 1995. "Towards a More Precise Decision Framework: A Separation of the Negative Utility of Chance from Diminishing Marginal Utility and the Preference for Safety," Theory and Decision 39:3 (November), 241-265.
- Powell, Melanie, and David Ansic. 1997. "Gender differences in risk behaviour in financial decision-making: An experimental analysis," Journal of Economic Psychology 18:6 (November), 605-628.
- Reinganum, M. 1981. "Misspecification of capital asset pricing: Empirical anomalies based on earnings yields and market values," *Journal of Financial Economics* 9 (1981), pp. 19-46.
- Royden, Halsey, and Patrick Fitzpatrick. 2010. Real Analysis (Fourth Edition). Prentice-Hall.
- Samuelson, P.A. 1938. "A Note on the Pure Theory of Consumer's Behavior," Economica, New Series, Vol. 5, No. 17, 61-71.

- Schmidt-Nielsen, Knut. 1997. Animal Physiology: Adaptation and Environment. Cambridge University Press.
- Schubert, Renate, Martin Brown, Matthias Gysler, and Hans Wolfgang Brachinger. 1999. "Financial Decision-Making: Are Women Really More Risk-Averse?" The American Economic Review, Vol. 89 No. 2 (May): 381-385.
- Sharpe, William F. 1964. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." Journal of Finance 19 (September), 425-442.
- Sillers, Donald A. 1980. "Measuring Risk Preferences of Rice Farmers in Neuva Ecija, the Philippines: An Experimental Approach," Ph.D. Dissertation, Department of Economics, Yale University.
- Stigler, George J. 1984. "Economics—The Imperial Science?" The Scandinavian Journal of Economics Vol. 86, No. 3 (September): 301-313.
- Stigler, George J. and Gary S. Becker. 1977. "De Gustibus Non Est Disputandum," The American Economic Review Vol. 67, No. 2 (Mar., 1977), pp. 76-90.
- Swiss Re. 2010. Sigma No. 2/2010: World Insurance in 2009. Zurich: Swiss Reinsurance Company Ltd. (http://media.swissre.com/documents/sigma2_2010_en.pdf).
- Yook, Ken C. and Robert Everett. 2003. Journal of Financial Planning (August)
- Wagenaar, W., 1988. Paradoxes of Gambling Behaviour. London: Erlbaum.
- Walker, Michael B. 1992. The Psychology of Gambling. International series in experimental social psychology. Elmsford, NY: Pergamon Press.
- Walker, Thomas S. 1980. *Decision making by farmers and by the national agricultural research program on the adoption and development of maize varieties in El Salvador*. Stanford: Food Research Institute.
- Weber, Elke U., A. Blais, and N. E. Betz. 2002. "A Domain-specific Risk-Attitude Scale: Measuring Risk Perceptions and Risk Behaviors," Journal of Behavioral Decision Making, 15 (August): 263-290.
- Zinkhan, G. and K. Karande. 1991. "Cultural and gender differences in risk taking behaviour among American and Spanish decision makers," Journal of Social Psychology 131:5, 741-742.

Figure 1: Portfolio Mean Returns versus Their Market Risk

(Prepared by authors from data in Black, 1993, Exhibits 3 and 4)

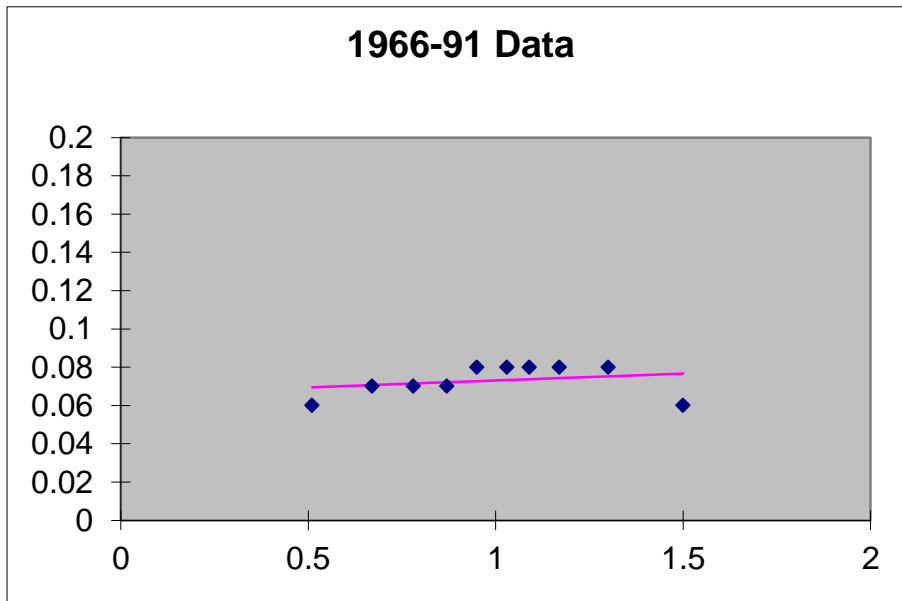
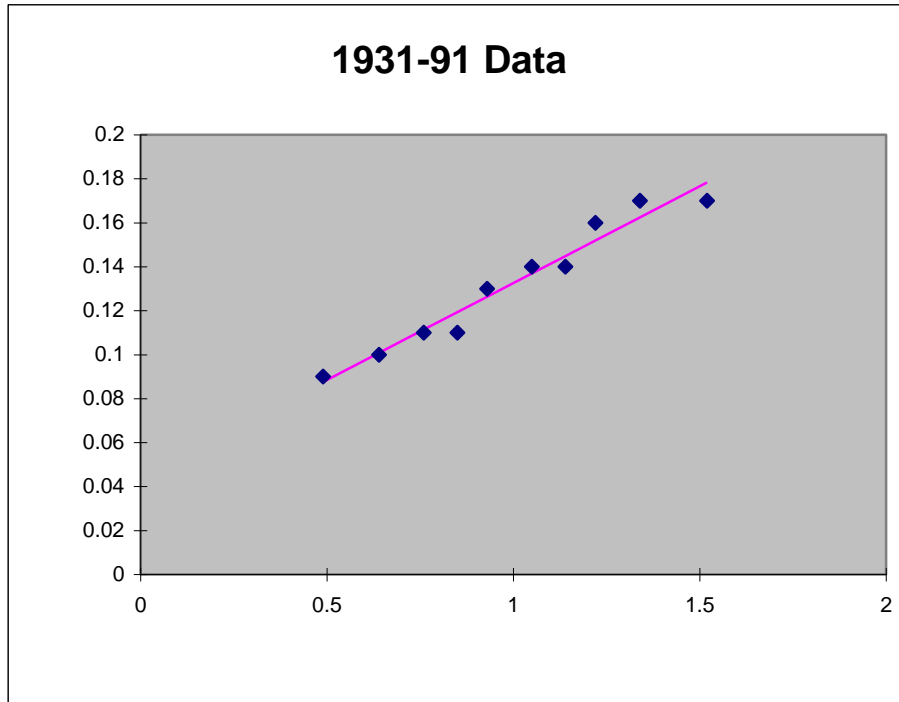


Figure 2: Net Payoff Functions ($y =$ net payoffs; $x =$ gross payoffs)

