Risky Curves: From Unobservable Utility to Observable Opportunity Sets

Daniel Friedman
University of California, Santa Cruz

Shyam Sunder
Yale University

Abstract

Fifty years of intensive search have yet to identify scientifically useful non-linear utility functions. Even complex functions with many free parameters do a poor job of predicting individual behavior in out-of-sample data, in new tasks and in new contexts. At the population level, socioeconomic and demographic data exhibit little power to explain variations in the curvature of estimated utility functions (risk preferences). There is no consensus even on the cross-sectional distributions of risk preference parameters. Qualitative accounts of macro-economic phenomena, financial markets, insurance, and gambling, the traditional justifications for non-linear utility functions, have not led to quantitative work of practical value.

The absence of scientifically useful evidence on curved utility functions suggests the merits of returning to the roots of choice theory. Placing the explanatory burden on potentially observable opportunity sets offers a simpler and more robust approach to understanding behavior under risk. Analysis of net payoff opportunities, including embedded options and other interactions with existing obligations, permits a parsimonious analysis of risky choice using only expected value, i.e., a linear utility function.

We are grateful to Rajneesh Mehra, Brian Mittendorf, and Steven Morris for conversations on the topic though we alone are responsible for the paper. Please do not quote without permission. Please send comments to dan@ucsc.edu or shyam.sunder@yale.edu. The is a preliminary version for presentation at CEBR Conference on Measuring Risk and Time Preferences (http://www.cebr.dk/researchtopmenu/events%20submenu/2004/spring/conference%202004.aspx).

Risky Curves: From Unobservable Utility to Observable Opportunity Sets

Daniel Friedman and Shyam Sunder

1. Introduction

Utility functions are neither deduced from fundamental propositions nor observed directly; they are inferred from observed behavior. The scientific value of such inferred functions derives from their usefulness in organizing and predicting out-of-sample observations. In this paper we examine utility functions inferred from observed choices under risk, and inquire if the inferred preferences exhibit sufficient stability and consistency to qualify as constructs of scientific value. We find little evidence so far that non-linear utility functions, despite their great flexibility and free parameters, provide an empirical advantage over the linear utility (or expected value) model of risky choice.

The expected value (EV) model is simple and widely applicable. It has no free parameters and is based only on the axioms of probability. Not surprisingly, its power to explain choice data falls short of perfection. Starting with Daniel Bernoulli's (1738) logarithmic or declining marginal utility through Kahneman and Tversky's (1979) prospect theory and its recent variants, scores of alternative models have been proposed. The shape of these alternative (non-linear) utility functions is defined by a set of free parameters, to be estimated from the observations the model is intended to explain. The model's scientific value depends on our ability to estimate the parameters. It becomes a useless tautology unless the parameter estimates have a degree of consistency and stability, or follow some knowable rule governing variations across contexts and time.

If the utility functions and their parameters were universal (applicable across persons, time, contexts and space), their estimation would use up only as many degrees of freedom as the number of parameters. Since the number of observations on choices under uncertainty is large, the number of degrees of freedom used up for estimation would be relatively small. We might infer, for example, that human beings act under uncertainty as if they have a constant absolute risk aversion (CARA class) utility function with parameter 0.5, and that the explanatory power of this model of universal human behavior exceeds the explanatory power of EV by, for example, 5 percent. The boost in explanatory power would easily justify the single free parameter and so this model would dominate EV. To the best of our knowledge, however, no such evidence has yet been adduced.

A second possibility is that the utility function or its parameters are specific to individual agents. Again, it should be possible to develop reliable estimates of the parameters from observations on the agent's risky choices. Even if the individual-specific parameters were totally idiosyncratic, each agent makes numerous risky choices; the number of parameters would have to be exceptionally large for us to worry about insufficient degrees of freedom. If the individual parameters were a function of some identifiable individual characteristics such as wealth, age, gender, social status, etc., we could estimate these functions from the observations by losing even fewer degrees of freedom. Attempts to identify such functions do not seem to have been successful.

The expected utility (EU) model and other alternatives to EV are often postulated to have their parameters vary not only by person, but also by time as well as context. There are no known constraints on how the parameters vary by observable characteristics of individuals (height, weight, age, wealth, etc.): education, contexts (buying food, clothing, or making investments, etc.); geography (Asia, villages, ski slopes, etc.); or time (night, sixteenth century, driving to office, etc.). There are not even any known constraints on how the statistical distributions of the parameters of these models across individuals, contexts, geography, or time are supposed vary, or remain invariant.

Does the presumed improvement in the explanatory power of the EU and the other alternatives to EV arise from something other than the advantage conferred by these free parameters? Not even Bill Gates can be richer than the man with a machine to print currency notes at no cost. Occam's Razor, favoring simplicity and economy, is science's defense against overly complex models with numerous free parameters. If the parameter estimates, and even their distributions, remain unidentified after a reasonable opportunity and effort, it becomes difficult to distinguish the model from a non-falsifiable matter of faith.

It is possible that the evidence not yet found may still turn up one day. In the meantime, our suggestion is to supplement the parsimonious EV model with a more careful examination of opportunities and constraints. As an illustration, consider two portfolios. Portfolio X consists entirely of US government bonds and insured certificates of deposit. Using the EU model, we might infer that Portfolio X belongs to a highly risk-averse individual with a concave utility function. Portfolio Y consists only of deep out-of-the-money call options on commodities futures. We might infer that its owner likes to take risk, i.e., has a convex utility function.

The range of permissible non-linear preferences in EU and its recent variants is so broad as to account for almost any possible behavior. The problem is that such easy explanations often lack predictive power. Even worse, they distract us from pressing the inquiry, and from possibly gaining deeper insights into how potentially observable circumstances may shape the observed behavior.

Portfolios X and Y in the above example happen to be managed by the same person—a financially astute colleague. She holds portfolio X for her 85 year old great uncle, a retiree who needs a steady income to maintain his present living situation. An unexpected increase in value would add to his bequest but hardly affect his lifestyle; an unexpected decrease could deprive him of quality care in his last years.

She holds portfolio Y for a national contest with a prize for the highest rate of return over the next 6 months. Winning the contest would give her career (and current wealth) a real boost, while coming up empty would have negligible impact. Hence her objective is not to maximize terminal portfolio value w, nor the expected value of some concave utility function $u(w)$. Rather, the objective is to maximize the probability of having the highest rate of return $r = (w-w_o)/w_o$ among all contestants. To maximize that probability, she wants and chooses a volatile portfolio, so the distribution of r has a fat upper tail.

In this example, fitting utility functions to observed choices provides no insights, nor has any predictive value. For example, it would fail to predict the impact of changing the rules of the contest, say to penalize volatility, and it would incorrectly predict our young colleague's behavior in other circumstances, e.g., in entering the state lottery or buying a home. Nor does the fitted utility function explain the portfolio choice in any meaningful way.

This example is contrived but we shall argue that the point is quite general. Section 2 of the paper critically examines existing studies that infer individual risk preference from observed behavior, and finds parameter inconsistency and instability is generic. Section 3 examines the evidence from larger scale economic phenomena of investments, insurance, and gambling to see what it reveals about attitudes toward risk in the population. Section 4 revisits the basic decision theoretic approach of economics, and argues that a careful analysis of all consequences of a decision, combined with EV criterion, can explain existing evidence better than EU and its more complex alternatives. Section 5 offers a discussion and some concluding remarks.

1A. Concepts of Risk

> *Risk: (n)(ca. 1661) 1. possibility of loss or injury: PERIL; 2. a dangerous element or factor; 3. the chance of loss. Synonyms: DANGER, hazard, jeopardy, peril. (vt)(1687) 1. to expose to hazard or danger (risked her life); 2. to incur the risk or danger of (risked breaking his neck). Synonyms: (1) ENDANGER, compromise, hazard, imperil, jeopardize, jeopardy, menace, peril; (2) VENTURE, adventure, chance, hazard, wager; (3) GAMBLE, chance, hazard, venture.*
>
> *To be "at risk"; to "take risks"; to be a "risktaker"; to have a "risk factor"—what these uses of the word risk have in common is a dependence on concepts of probability and harm. However, there the similarity ends. Intuitively, we know that being at risk for a harmful outcome does not necessarily require risk-taking, and that a person with a risk factor may not take risks. Within disciplines, these terms have their own precise definitions; thus epidemiologists know what they mean by relative risk, industrial engineers know what risk assessment is and how to do it, and psychologists construct psychometric scales measuring risk-taking propensity and risk perception. Definitional problems may arise, however, when disciplines collide. Leigh (1999).*

The common meaning of risk implies potential harm and uncertainty. Exposure to risk is something to be assessed, managed, controlled, mitigated, or avoided altogether. This is the sense in which risk is referred to in insurance, health and medicine, corporate governance and internal controls, engineering, human development, credit and banking. Other things being the same, few would knowingly want to choose greater exposure to risk of a loss. It follows directly from the definition of loss as something undesirable. Words liability, indemnity, compliance, warranties and guarantees and assurance often accompany risk in various contexts.

Risk is recognized to be a multidimensional concept (Carney 1971, Lipsitt and Mitnick 1991). While it is negative in most contexts, a positive connotation also exists as in adventure, bravery, heroics, and entrepreneurship. Risk has a time dimension, being conceptualized as short-term (acute illness) or long-term (heart disease), and instantaneous (accident) or cumulative (smoking). Risk taking has been viewed as a generalized tendency of individuals which holds across contexts and types of risks. It has also been visualized as being specific to domains such as monetary, physical, ethical and social (Jackson et al. 1972). Any tendency toward risk-taking in such specific domains does not carry over into other domains (MacCrimmon and Wehrung 1986). Knight (1921) distinguished between risk and uncertainty on the basis of known probabilities in the former, and unknown in the later. In certain fields such as public health, the distinction between individuals' subjective perception of risk and population level objective risk is important (Jeffrey, 1989).

In the banking industry, one refers to various elements of risk such as credit, liquidity, settlement, market, and legal. Credit risk, for example, is associated with counterparty to a transaction not fulfilling its obligations in the transaction in full by the due date.[1] All these dimensions of risk refer to the possibility of harm.

---

[1] "Credit risk is most simply defined as the potential that a bank borrower or counterparty will fail to meet its obligations in accordance with agreed terms." Bank of International Settlements (1999).

In spite of the predominance of association of harm and uncertainty with most concepts of risk, dispersion of outcomes,[frequently operationalized as variance (or standard deviation)], now dominates many treatments of risk in economics, especially financial economics. Dispersion of outcomes from the mean captures uncertainty; but it gives the same weight to positive and negative deviations from the mean, and does not focus on the "harm" part of the concept of risk. Markowitz (1952) recognized that as a measure of risk variance, which has powerful arguments in its favor. Mathematical tractability of this concept has helped create the modern portfolio and capital asset pricing theories. Statistical and econometric convenience of variance has helped build a long empirical tradition in financial economics.

This convenience of the dispersion-based concept of risk tends to obscure its radically different nature that does not capture the sense in which the term risk is used in most domains of human endeavor, including financial domains such as insurance and banking. After discussing risk in terms of its variance-based concept in the following sections, we shall return to its harm-and-uncertainty conceptualization in the final section of the paper.

2. Inferring Individual Risk Preferences from Observed Choice

When utility function is linear, the expected utility (EU) model is equivalent to the expected value (EV) model. Since the purpose of the paper is to inquire into the scientific usefulness of postulating curved utility functions, it is convenient to present the discussion in terms of EU model.

EU applies to a decision maker (DM) choosing among risky alternatives $A = \{(x,p),\ (y,q), (z,w), \ldots\}$, where each alternative (or gamble) has a known vector x of possible outcomes, e.g., with components $x_1 < \ldots < x_n$, and corresponding known vector p of probabilities, e.g., $p_1, \ldots, p_n \geq 0$ such that $\sum p_i = 1$. According to the EU model, the DM has some continuous, strictly increasing utility function u: $\boldsymbol{R} \rightarrow \boldsymbol{R}$, normalized so that $u(0) = 0$ and $u(1) = 1$, such that she chooses the alternative in A with highest expected utility, e.g., $\mathrm{E}u(x) = \sum p_i\, u(x_i)$.

Standard definitions associated with the EU model will facilitate later discussion. The certain-equivalent of a gamble (x,p) is the outcome c such that $\mathrm{E}u(x) = u(c)$. The expected value of a gamble (x,p) is $\mathrm{E}V = \sum p_i\, x_i$. When u is concave, e.g., $u'' < 0$, then the certain-equivalent is less than the EV for every non-trivial gamble, and the person is said to be risk averse. The opposite is true when u is convex, e.g., $u'' > 0$, in which case the person is said to be risk seeking.

The EU literature deals with many issues, some of them quite complex and controversial. For example:

- Can du/dx be interpreted in a risk free environment as the marginal utility of money (e.g., Friedman and Savage, 1952)?
- Can the outcomes x be interpreted as wealth increments or only as terminal wealth levels (e.g., Cox and Sidiraj, 2003; Rabin, 2000)?
- Should the EU be generalized or supplemented by the prospect theory (Kahneman and Tversky, 1979), the dual theory (Yaari, 1987), or one of numerous other proposals?

We touch on these issues in later sections, but our focus in this section is a simple empirical question: is the utility function u an intrinsic personal characteristic? That is, does a specific person tend to choose, with some consistency, among risky alternatives as if maximizing the expectation of a specific utility function u?

The axiomatic derivation of the EU model assumes an affirmative answer. Classic treatments (e.g., von Neumann and Morgenstern, 1944; Arrow, 1971) and recent textbook treatments (e.g., Mas-Colell et al., 2000) show that if risky choices satisfy a short list of plausible consistency axioms, then there is a particular utility function whose expectation is always maximized. The utility function is defined up to a positive affine transformation, so the normalization $u(0) = 0$ and $u(1) = 1$ produces a unique utility function. That is, any decision maker who is consistent in the sense of the axioms (i.e., the choices satisfy very sensible conditions such as continuity and independence) can be identified with a particular element u in U, the infinite-dimensional family of admissible utility functions.[2]

The theory has a straightforward empirical interpretation: observed choices in a sufficient variety of risky situations allow us to infer (or closely approximate) a person's utility function $u \varepsilon U$. Further, this function can be used to predict that person's choices in novel risky situations. The first step—estimation of u from choice data—is a mechanical process which always yields an output. Whether this estimate has a predictive value for out-of-sample data is an empirical question.

Universal utility functions.

The strongest empirical interpretation of the EU model is that a single utility function u works for everybody, capturing most choices of most people. This function might be linear, or a member of some concave family such as CARA or CRRA,[3] or perhaps something more complicated. Markowitz (1952) and Friedman and Savage (1948), for example, proposed universal functions that have concave as well as convex segments.

This interpretation was soon seen to be too strong to find support in data. Empirical studies beginning in the 1950s found systematic interpersonal differences in choices (e.g., Edwards 1953, 1955). The first fallback position is that although not everyone has identical u, every individual belongs to one of just a few basic risk types, analogous to blood types. Once a person has been identified to have Type O, A, or B blood, the identification is good forever and in all circumstances. Unfortunately, as detailed below, no such types of utility functions have yet been identified with any acceptable degree of reliability or stability.

Systematic variation in utility functions

The next possible interpretation is that observable personal characteristics map nicely into a particular utility function. That is, a person's utility function has some simple and knowable relation to his/her age, gender, wealth, education, etc., so economists can make useful empirical generalizations of the following form:

- Lower middle class American males of age 30 typically have a Friedman-Savage type utility function with lower inflection point near income 0 and upper inflection point near 20, and absolute risk aversion coefficients in the three segments of approximately $a = 2.5$, -1.2, and 2.5 respectively.
- An upper middle income Japanese housewife of age 50 is likely to have a CRRA utility function with parameter $r \approx 3.0$.

The prospect of making such statements seems to have inspired many empirical studies, but the evidence has not been supportive. The most favorable and heavily cited study of this kind is a field experiment with more than 100 farmers in India, described in Binswanger (1980, 1981, 1982). The experimental task was to choose one of eight alternative bets of form (x_1, x_2) with $0 < x_1 \leq x_2$ and $p_1 = p_2 = ½$. One extreme alternative had no risk with $x_1 = x_2 = 50$ points, and the

[2] A technical clarification: The Stone-Weierstrass theorem ensures that we can approximate the continuous increasing utility function arbitrarily closely over a bounded interval by a twice continuously differentiable function. The space U of such smooth utility functions has a countably infinite basis; see e.g., Royden (1968).

[3] A utility function u has constant absolute risk aversion (CARA) if the function a(x) = -u''(x)/u'(x) is constant, i.e., independent of x. It has constant relative risk aversion (CRRA) if the function r(x) = -xu''(x)/u'(x) is constant in x. Textbook treatments show that unnormalized CARA utility functions have the form …

other extreme alternative was $x_1 = 0$, $x_2 = 200$. The intermediate alternatives were chosen so that the risk (here proportional to the 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.[4] More germane to the present discussion, he estimated the impact of demographic characteristics on a risk aversion parameter intended to summarize an individual's utility function (1981, Table 2). He found that wealth, schooling, age, and caste were all insignificant; the only variable reported with significant impact was Luck![5] Sillers (1980) found that the risk attitudes estimated for Filipino farmers in his study showed risk aversion similar to Binswanger's Indian farmers, especially when the stakes were high. However, despite considerable effort, neither Binswanger nor Sillers were able to account for the individual farmers' decisions involving familiar risks over alternative crops, fertilizer use, etc., using the risk preferences inferred from the experiment.[6]

In a recent study, Harrison et al. (2004) conducted a field experiment to estimate the risk attitudes in Denmark. They conclude that the estimated risk attitudes show significant deviation from risk neutrality for large but not small amounts, vary across identifiable populations, and relative risk aversion tends to rise as lottery amounts in risky choices increase. Unlike Binswanger and Sillers studies, Harrison et al. do not report an attempt to validate the risk attitudes estimated from their field experiment with lottery choices in out-of-sample domains of choice.

Our search of the empirical literature revealed no consensus on how *any* potentially observable personal characteristic affects the utility function. Studies employed a variety of methods, including surveys and laboratory experiments, as well as the occasional econometric analysis of panel data. Surely the strongest candidate for a consensus view concerns gender. Responses to survey questionnaires consistently indicate that women on average perceive greater risk than men in a variety of activities, personal, as well as social and environmental, 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. However, neither of these findings speaks directly to whether utility functions differ by gender: the survey data differences may arise mainly from informational (or response bias) differences (cf., Weber, Blais and Betz, 2002), and arrest record differences may 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 effects of risk preferences. There have been dozens of relevant studies, and many of them corroborate the conventional view that

---

[4] The conclusion is arguable. Binswanger conflates stake size with experience and time (and perhaps wealth), and the highest stakes condition is hypothetical. The farmers' modal choices are always at the adjacent bets B=(40, 120) or C=(30, 150), with EVs 80 and 90. As stake size increases, the farmers' modal choice moves from B to C and (at highest stakes) back to B, and the distribution tightens around the mode. Binswanger doesn't report any tests of the null hypothesis that the mean or median choice changes with stakes; it appears that most such tests would fail to reject. On the other hand, the conclusion is supported in a recent laboratory experiment reported in Holt and Laury (2002).

[5] 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.

[6] "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, apparently because of measurement errors and the small size of the sample." Sillers (1980, p. 211).

women tend to be more risk averse than men. Powell and Ansic (1997), for example, report that their female subjects were less risk *seeking* in laboratory tasks than the males. 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. If the shape of the utility function depends on the nature of the task, and the task is not included as a parameter in the nonlinear utility theories, the scientific value of the construct itself becomes doubtful.[7]

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 their under-weighing the low probability events, and overweighing the high probability events. This tendency diminishes with age and disappears among adults. Otherwise age has little discernable impact on risk preferences. We found some scattered results regarding ethnicity. For example, Zhinkhan et al. (1991) found that the Spanish subjects were more willing to take risks than Americans—but there is no real consensus. Harrison et al. (2003) report on a field experiment in Denmark that showed no age or gender effects but indicated an education effect. 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. Income variable loaded significantly, but may have been confounded with how investment companies assign risk tolerance scores.[8]

Finally, 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 found large performance differences in one complicated risky task, and the standard interpretation is that the VM brain structures are involved in making risky choices. However, the authors found no significant differences between the two groups for any of their simple risky tasks, and cite other studies that have 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.

---

[7] See Harrison et al. (2003, footnote 25) for a summary of other studies of differentiation of risk attitudes by gender. These include Levin et al. (1988), Powell and Ansic (1997), Harbaugh et al. (2002), Holt and Laury (2002), Brown et al. (1999), Jianakoplos and Bernasek (1998), Rubin and Paul (1979), and Robson (1996).

[8] 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 is highly likely to be a function of income and wealth.

Our partial scanning of the literature has not yet revealed validated theory or consistent evidence to systematically link the curvature of utility functions to some observable characteristics of human beings.

Idiosyncratic utility functions.

Although socioeconomic, demographic, and even medical information may not tell us much about a person's risk attitudes, there still may be intrinsic utility functions. Perhaps one can identify some overall distribution of utility functions across the population. Knowing such a distribution would have empirical value—Goldman Sachs or Merrill Lynch, for example, could use tables of the distribution to hone their financial products and market them more efficiently. Unfortunately we have not found any studies reporting such distributions, nor are we aware of any such proprietary risk preference tables used in industry.[9]

A weak version of intrinsic utility is that each person's utility function is an arbitrary element of *U*. The elements may be so scattered and heterogeneous that useful distributions can't be tabulated, but at least a person's intrinsic *u* remains the same across choice tasks. In this version, we have no a priori predictions regarding risky choice until we have estimated each person's utility function from a sufficient variety of past choices.

Available evidence is unkind even to this mild hypothesis: the utility function seems to change drastically as we change the way it is measured.[10] The most favorable study we could find was Harlow and Brown (1990), who measured risk attitudes of more than 100 subjects in four distinct ways: (a) from bids in first price auctions with independent private values, using a particular bidding model known as CRRAM, (b) from responses to the widely used MMPI survey questions, (c) from responses to another psychometric survey, SSSV, and (d) from 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,[11] but no relation for the women subjects.

A short discussion of risk aversion in laboratory auctions may be useful because it has been investigated intensively. First-price sealed-bid independent-private-values auctions have a risk neutral bidding formula, and observed bids in lab experiments typically are higher. Risk aversion is the most popular explanation, as worked out most carefully in CRRAM (Cox et al., 1988). Selten and Ockenfels (2002) challenge this explanation, and show that the steady state of a plausible adaptive process ("impulse balance equilibrium") can also explain overbidding as well as the effect of information treatments. Risk aversion can account for the information effects only if risk preferences are not intrinsic, but instead change drastically as more or less information becomes available.

In theory, risk aversion leads to lower bids than 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. Steven Kachelmeier and Mohamed Shehata (1992) infer

[9] Industry analysts routinely assess the size of the existing or potential markets by consulting tables that describe the distribution of people who hold securities or insurance policies of various kinds. But nonlinear utility functions do not yet play any part that we know of in the conceptualization, construction, or use of such tables.

[10] This is even true for survey data. Weber, Blais and Betz (2002) compare people's responses to survey questions across five content domains (investing, gambling, health/safety, recreation, ethics and social). There is surprisingly little correlation across domains, typically 0 to 0.6, despite the intention of the study to find an integrated measure. The focus in the text, however, is data from laboratory experiments with salient rewards.

[11] Even this positive result is undercut by the fact that estimates of the two other CRRAM parameters (with no theoretical relation to risk) have correlations with (b)-(d) of about the same significance.

risk-seeking preferences when their subjects sell a gamble, and infer risk-averse preferences when the same subjects buy the gamble. The Becker-DeGroot-Marschak mechanism is the standard tool for eliciting a buying price for a gamble. Several studies, including Berg and Reitz (1997) and Berg, Dickhaut and McCabe (1992), and Fong and McCabe (1999, PNAS) find that the nature of (and personal involvement in) the task pushes measured risk aversion up or down.

A generous interpretation of the last several studies is that intrinsic risk attitudes exist, but are overlaid by task-specific measurement biases. This interpretation is hard to refute when comparing risk measurements across different sets of subjects, but it is put to the test when the same subject is measured in several different ways. The Harlow and Brown (1990) results were not encouraging, although only one of their measurements had direct implications for utility functions, and the study leaves some room for hope that it might be possible to obtain data to support the concept of risk attitude as an intrinsic characteristic of individuals.

Isaac and James (2000) close the door on the last hope for intrinsic utility functions. They found a strong *negative* correlation across individuals between risk aversion as measured in a first price auction, and as measured in a standard (Becker-DeGroot-Marschak) procedure. The separate measurements corroborate earlier studies, so we can conclude that no additive bias correction, nor even a monotone transformation, can account for the inconsistent measurements across tasks.

In financial planning literature, which is concerned with individual investment advice customized for investment clients, their risk attitude, capacity, knowledge and propensity are said to be relevant considerations. Cordell (2001) postulates complex mutual dependencies among the four variables with age, income and wealth explicit determinants of risk capacity.

Even if u is not intrinsic to an individual, the EU model still might be empirically useful if u responds to the context according to some knowable laws. For example, utility of the same person may take one form in choosing jobs, another in choosing financial investments, and yet another during sports or adventure. This context dependence itself could vary by individual according to some knowable patterns. However, we were not able to find any empirical evidence supporting the discovery of such laws, either at the individual level or in aggregate.

A final possibility is yet another hypothesis—the weakest possible, in our opinion—that utility functions are specific to each individuals in each of many possible contexts, such as investment, insurance, gambling, health, sports, etc. Thus in a world of N individuals and M contexts, we should allow for, and estimate, a total of $N \bullet M$ utility functions. Given the number of actions people take, there should be enough degrees of freedom in the collectible data to place such estimation within the realm of possibility. However, we are not aware of any such efforts that have yielded results with useful out-of-sample predictive power.

The conclusion is unpalatable but clear: the non-linear utility model has not yet demonstrated any empirical value at the individual level. If measured utility u shifts in arbitrary ways across individuals and tasks, and therefore the model has no more predictive value other than attributing choices to themselves at best, or to arbitrary, unobservable moods or spirits—human or animal—placing us well outside the domain of science.

3. Using Risk Preferences to Explain Macro Observations

Even if it is not possible to infer individual risk preferences from micro-level behavior, the nonlinear utility models might nevertheless yield distinctive insights into important economy-wide phenomena such as securities markets, insurance, and gambling. Let us examine the evidence.

Stock Markets

Risk aversion is the crucial ingredient in prominent theories of investor behavior. Markowitz' (1952) portfolio theory is built on the assumption that investors consider the mean and the variance of returns, and no other characteristics, when they choose a portfolio of investments. This assumption follows from the EU model in important special cases, e.g., quadratic *u*, or multivariate normal securities returns. The equilibrium consequence of this portfolio theory is that the mean return of each security *I* is an increasing linear function of its market risk β = Cov (*I, M*)/Var (*M*), the ratio of covariance of the security and market returns to the variance of market returns (Sharpe 1964 and Lintner 1965). This theoretical linkage between the mean returns and the market risk of investments is often mentioned as a foundation stone of the modern investment theory.

A positive relationship between mean returns and market risk has been the subject of many careful econometric analyses. Empirical support for the proposition is questionable at best. Fischer Black and Eugene Fama are two of the founders of the theory of risk-return relationship. After many years of careful econometric scrutiny of US stock market data, Fama and French conclude that it does not support theoretical linkage between risk and return.

> 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 βs. (Fama and French, 1992, p. 428).

The business press has reacted to these findings with articles such as "Is Beta Dead?" (Wallace 1980). Fischer (1993) reported a positive relationship for the period 1931-91 but virtually no relationship during 1966-91; the mean returns on portfolios of varying market risk were about the same during the latter 25-year period (see Exhibit ??). However, he disputes the "beta is dead" conclusion. Kothari, Shanken, and Sloan (1992) conclude that beta plays a significant role in determining equity returns, allowing room for other determinants also. A recent edition of a finance textbook summarizes the current state of affairs as follows:

> Since William Sharpe published his seminal paper on CAPM (*capital asset pricing model*), researchers have subjected the model to numerous empirical tests. Early on, most of these tests seem to support the CAPM's main predictions. Over time, however, evidence mounted indicating that the CAPM had serious flaws (Smart, Magginson and Gitman 2004, pp. 210-212).

The textbook goes on to list the difficulties of properly testing (read *finding evidence in favor of*) the CAPM model, the difficulties hardly mentioned to question the early evidence that was thought to be supportive of the CAPM. These difficulties include: (1) unobservability of the dependent variable of the model—the expected return of securities; (2) unobservability of the expectation of market risk and the errors inherent in projecting risk estimated from the past data—the longer the data series used to estimate market risk, larger the likely error from projection; (3) the difficulty of identifying the correct risk free rate of return; and (4) uncertainty about the market portfolio and of the expected return on the market portfolio.[12] Conspicuous by its absence from the list is the possibility that assumptions regarding risk aversion may be invalid.

Brealey and Myers (1996), one of the best-known textbooks in finance, after reviewing the literature, shift the burden of evidence 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).

[12] Roll (1977) is a good review of this the problems of testing the theory.

Chan and Lokanishok (1993) "do not feel that the evidence for discarding beta is clear-cut and overwhelming." This is hardly a ringing endorsement of the idea that stock market phenomena at are consistent with the investors at large having concave utility functions.

If the behavior of the U.S. stock market, arguably the most liquid and certainly the most prominent market in the world cannot be explained on the basis of non-linear utility functions, is it possible to explain investor choice without assuming some form of risk aversion? Risk can be measured in many different ways; variance is just one special case. Different levels and kinds of risk change the opportunity sets available to agents. Analysis of the effects of risks on opportunity sets and choices might yield better insight into asset pricing puzzles than postulating unobservable utilities. Indeed, options theory takes just such an approach and has enjoyed considerable success in explaining securities market behavior in the last 30 years.[13] In the next section we will return to the relation between risk and opportunity sets.

Bond Markets

Risk analysis is different in the bond market. The popular ratings of bond risk—by Standard & Poor's and Moody's—are not based on the dispersion of outcomes that lead agents with concave utility functions to demand risk premiums. Instead, the ratings 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.[14] 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. Higher yields to maturity associated with lower-rate bonds cannot, therefore, be taken as evidence that bondholders have concave utility functions.

Fisher (1959) postulated that corporate bond yields in excess of yields on treasury securities are determined by two factors: chances of default and marketability of bonds. He postulated the chances of default to be a function of earnings variability, time elapsed since previous default and leverage, and the marketability to be a measured by the value of publicly traded bonds outstanding. His statistical regressions in log form for five cross sectional samples taken from 1927, 1932, 1937, 1949 and 1953 explained 74 percent of the variability of excess yields for 366 firms in terms of these four variables. Concavity of the utility function was not considered.

Altman ([1989], p. 918, Table 5 and Figure 1) examined the promised return on bonds of all ratings net of actual defaults. He found that the link between the bond ratings and this net yield is monotone increasing, except for the two riskiest categories (B and CCC) where the link breaks three years after the bonds are issued. The excess yields on bonds generally increase for lower rated bonds

The monotonic link is a puzzle from a CAPM perspective: defaults are mainly idiosyncratic so the betas and risk premiums should be zero. Possible explanations for the

---

[13] "What would I specialize in if I was starting over and entering the field today? … I reduce my advice to a single word: options." Miller (1999).

[14] For example, Moody's Aaa and Baa ratings are described as follows: Bonds which are rated Aaa are judged to be of the best quality. They carry the smallest degree of investment risk and are generally referred to as "gilt edged." Interest payments are protected by a large or an exceptionally stable margin and principal is secure. While the various protective elements are likely to change, such changes as can be visualized are most unlikely to impair the fundamentally strong position of such issues. Bonds which are rated Ba are judged to have speculative elements; their future cannot be considered as well-assured. Often the protection of interest and principal payments may be very moderate, and thereby not well safeguarded during both good and bad times over the future. Uncertainty of position characterizes bonds in this class.

observed link include (a) investors hold undiversified bond portfolios, (b) defaults are correlated with the market returns, (c) market inefficiency, e.g., demand limitations due to risk class investment restrictions, (d) the value placed on liquidity, and (e) interactions with interest rate risk. Explanations (a) and (b) require the additional assumptions that bond ratings are related to dispersion measures of risk, and that investors have concave utility functions. The other explanations do not require assumptions regarding utility functions. We do not yet know which of these explanations best fit the data. We are not able to locate evidence that supports the proposition that bond investors have concave instead of linear utility functions.

Insurance

In 2001, the insurance industry collected $2.4 trillion in premiums world-wide, and $904 billion in the US alone, accounting for a significant part of the economy. The negative actuarial value of insurance policies is often cited as evidence of widespread risk aversion. Even if we are not able to measure the degree and form of risk aversion of specific individuals or classes of individuals, the existence of such a large insurance industry must, the argument goes, convince the skeptics of the prevalence of concave utility functions over the outcomes of insured events.

We are not persuaded. The next section will propose an alternative explanation, but at this juncture a few remarks are in order. First, with insurance as with bonds, the relevant risk measure seems to be the probability of loss, as opposed to the dispersion of outcomes. In the dispersion-based perspective, risk averse decision makers dislike deviation of outcomes from the mean in both positive as well as negative directions. Few insurance policies aim primarily at the dispersion of outcomes in this sense. Instead, the contracts seem more analogous to option contracts, covering one tail of the outcome distribution—protecting the policyholders in case of a loss, but not penalizing them in case of a gain. (Option theory is ideally suited to analyze the value of such contracts to the policyholders; once we consider the option values, insurance contracts with negative actuarial value can be valuable even in absence of risk aversion. **[Analysis of insurance using options theory]).**

Second, insurance can be seen as a way of simplifying one's life by reducing the number and diversity of contingencies one must plan for, both in personal life as well as in organizations. In decision theoretic terms, it helps chop some branches off the decision tree that must be analyzed and planned for. An auto insurance policy matches the loss from an accident to the insurance reimbursement, allowing the policyholder to ignore that contingency after the premium has been paid. The life insurance policy matches the loss of income of the family breadwinner to the insurance proceeds, again allowing the family to reduce the worry about such a contingency. Indeed, this is the standard insurance sales pitch. Planning for contingencies is costly, both in personal life, as well as in managing organizations. Once the planning costs saved by insurance are taken into account, insurance policies may no longer have a negative expected value to the holder.

Third, it is worth noting that prospect theory, one of the better-known variations of nonlinear utility, suggests that people in general are risk averse to gains to their current wealth and risk loving toward losses. One implication of loss aversion is that people in general should refuse to buy insurance unless the actuarial value of the policies is sufficiently positive to compensate them for the opportunity to enjoy the risk they lose in buying insurance. With the possible exception of government subsidized flood insurance, we know of no insurance policy with zero, much less positive, actuarial value.[15]

Finally, insurance policies do not exist in primitive societies. They are contractual forms developed in certain civilizations. It is unclear whether there is some intrinsic desire to purchase insurance for home, car, life or health, or whether the desire arises instead from some social learning process, possibly abetted by the industry's marketing efforts.

Gambling

A recent study estimated that in the US alone, $550 billion annually is wagered in organized gambling (National Research Council, 1999). Some eighty percent of US adults report having engaged in gambling at some time in their lives, and a significant minority are heavy gamblers. Though not as large as security markets and insurance, gambling is surely large and pervasive enough to deserve theoretical attention.

Just as economists invoke concave utility to explain insurance, they invoke convex utility to explain gambling. For a strictly convex utility function, the certainty-equivalent is larger than the expected value, and so some negative expected value gambles will be accepted. However, the strange implications of the model are seldom spelled out. Someone who maximizes the expectation of a globally convex utility will, at any fixed degree of actuary unfairness, prefer the largest possible gamble: a mean preserving spread always increases expected utility. A Markowitz-type utility function still predicts a preference for gambles with an infinite downside, because the convex domain has no lower bound. Preferred bets for a Friedman-Savage utility 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 fixed degree of unfairness) involves only two possible outcomes a and d such that $a<b<c<d$ and $u'(a) = u'(d)$.[16] Also contrary to common sense, the model predicts that 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 over the convex domain.

Gambling is big business and 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 utility function, however tortuously twisted, accounts only for the monetary consequences. It ignores the thrill, the testosterone, 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.

Even though the incidence of pathological gambling in US is less than two percent, we cannot discount the importance of non-monetary associations and consequences of gambling in seeking an understanding of gambling behavior. Traditional Freudian analyses of compulsive gambling point to masochistic self-punishment, autoerotic control, and even oedipal issues (Crance, [1988], p. 54). More recent approaches point to events outside of the gambler's

[15] We asked two behavioral economists how they reconcile prospect theory with the existence of the insurance industry. One argued that a different bias, overweighting of small probabilities, overcomes the impact of loss aversion. The other correspondent cited {I've lost my list of Jamal's defense}.

[16] The outcomes are points of tangency with the most northwesterly line that intersects the graph of the utility function; see Figure 6 below. For the FS interpretation to have empirical content (see Friedman 1953), we would have to either know the parameters of the utility function, or know the laws that determine the values of the parameters for individuals or contexts. To the best of our knowledge, neither of these conditions has been fulfilled after some five decades of intensive investigations on utility theory.

personal control and influence (Dickerson, 1984), a form of safe 'risk taking' to relieve tension (Crance, p. 58), and an "avenue of escape from routine and boredom" (Ezell, 1960).

Behavioral psychologists have their own explanations of gambling behavior, based in part on reinforcement learning theory. Lotteries in particular have been described as a "variable ratio" form of Skinnerian conditioning, offering reward at unpredictable intervals - resulting in the belief that winning is inevitable. (Chamberlain and Cowan, 2000, p. 69). This view is confirmed by studies that have shown that "state lotteries… have increasingly changed the structure of lotteries to take advantage of cognitive biases and responses to reward," changing lottery format from a single large prize with little possibility of winning to include also many small prizes (NRC, 1999). For example, in the U.S. lottery, a small fraction of the tickets pay a prize less than the ticket price. Supposedly, such "winners" are more likely to increase ticket purchases even though they just experienced a net loss (Wagenaar, p.70).

Extended discussions of the psychology of gambling can be found in Michael B. Walker (1992), Chris Gudgeon (1995) and http://www.chass.utoronto.ca/~johnbell/Final/possessionritual.html . The industry has found ways to differentiate the product to appeal to people in a wide range of age groups, socioeconomic classes, educational attainment, and gender. Conspicuous by its absence in the applied literature is any reference to or investigation of the curvature of utility functions.

4. Non-linear Opportunity Functions

When the evidence is unkind to a cherished hypothesis, a scientist's instinct is to go back to first principles. First principles in decision theory are that the decision maker (DM) chooses the most preferred available opportunity. Perhaps, then, it is time to take a closer look at the nature of risky opportunities and the content of preferences.

Expected utility theory (along with most variants) assumes that in risky choice situations, the available opportunities are lotteries whose prices are wealth increments, and that preferences are defined over wealth or wealth increments. However, DM preferences in standard choice theory are defined over final consumption bundles or flows, and wealth or wealth increments enter only the indirect utility function. Thus there is a gap between net payoff—what the DM truly values—and gross payoff—the lottery prizes.

We shall now argue that many of the phenomena rationalized by non-linear utility can also be explained by expressing net payoff as a non-linear function of gross payoff. Putting the non-linearity in the opportunities rather than the preferences has two crucial advantages. First, it is much more parsimonious: all parameters potentially are observable, rather than inferred. Second, it offers novel predictions: shifts in opportunities that affect the net payoff function will systematically affect conventional measurements of risk aversion.

Concave Net Payoffs.

We begin with circumstances for which net payoff is a concave function of gross payoff. Suppose the 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 ε (0, 1) of the shortfall. For example, if the DM has a credit card balance of z = \$1000 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 needed 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).

Figure 1 shows the resulting net payoff $n(g) = g-z$ for $g>z$ and $n(g) = (1+a)(g-z)$ for $g<z$. 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 $n(g)$ is strictly concave over the support of z.

Progressive income taxes have a similar impact: the slope of the function n is less at higher g due to higher marginal tax rates. Fiduciary responsibilities also lead to concave net payoffs for the trustee. When he obtains a gross payoff for the client far above the expectations, his net payoff is only slightly higher than when meeting expectations, but when the gross payoff falls short of expectations his net payoff is far lower after taking into account the costs of a poorer reputation and perhaps legal costs.

Discrete, irreversible decisions also lead to concave net payoff functions. For example, suppose you 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. You might conclude that this DM deducts a risk premium from the offer, and so you might infer that the DM has a concave utility function. However, the deduction might not be a risk premium. It could instead be the value of the wait option: favorable new job offers might be more likely for an established incumbent than a new hire in a new city. Dixit and Pindyck (1994), for example, show that the deducting the value of such options leads to net payoff that is concave in the job offer g.

In all these cases, an ignorant 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 will be able to make the distinction. He can avoid the specification error of attributing an unstable concave utility function to a risk-neutral DM with varying net payoff functions.

Convex Net Payoffs.

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

More important examples include decisions in the shadow of bankruptcy, or bailout. Suppose that failure to meet a contractual obligation $z>0$ results in a bankruptcy proceedings and shortfalls are passed to creditors, as in Figure 3. The net payoff again is $n(g) = g-z$ for $g>z$ but now is $n(g) = (1-a)(g-z)$ for $g<z$, where $a \in (0, 1)$ is the share of shortfall borne by other parties. Again a random component to cash flows smoothes out the graph and here makes n a strictly convex function over the support of the uncertainty.

Bailouts create convex net payoffs in exactly the same way. Perhaps the most famous example in the US is the incentives facing the savings & loan banks in the 1980s. 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 concave 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 deposit insurance.

More Complex Net Payoffs.

Certain opportunity sets would lead an ignorant outside observer of a risk-neutral DM to infer a non-linear utility function with concave and convex segments, as in Markowitz or Friedman-Savage. For example, suppose the DM lives in subsidized housing with subsidy rate $a > 0$ if her income is less than or equal to z, and becomes ineligible for subsidy if actual income (taking into account opportunities to disguise it) exceeds $y>z$. If ineligible, she spends fraction $c>0$ of incremental income on housing. Then net income (after housing) is $n(g) = n_o +(1-a)(g-z)$ for $g<z$, and $n(g) = n_o + (g-z)$ for $z \le g \le y$, and $n(g) = n_o+y-z +(1-c)(g-z)$ for $g > y$; see Figure 4. 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 utility function over gross income.

Friedman and Savage (1948) motivate their proposed utility function with a story about the possibility of the DM moving up a class on social ladder. To sharpen their story a bit, suppose that z 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 w one 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 $y>z$ 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, after the usual smoothing, one obtains the Friedman-Savage function in Figure 5 with inflection points in the intervals (w, z) and (z, y). But the characteristic non-linear shape reflects opportunities in the net payoff function $n(g)$, not some sort of intrinsic preferences.

Marshall (1984) obtains a similar shape for the indirect utility function for income. True preferences are assumed 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.

Nils H. Hakansson (1970) derives a Friedman-Savage type function in an additively-separable multiperiod setting. The net payoff is expected utility of wealth, given by a Bellman equation for the consumption-investment plan, assuming that the utility function of consumption each period is CRRA. The gross payoff is the present value of endowed income. Hakansson derives the desired utility function explicitly from particular constraints on investment and borrowing.

Robert T. 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 are the usual sort of substitutes (decreasing MRS). Masson shows that realistic capital market constraints can create concave or mixed functions $n(g)$, where g is first period endowment and n 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 $n' = b$ when realized g is sufficiently small, and will lend so $n' = l$ when realized g is sufficiently large. For intermediate values of g the DM consumes all of the incremental first period endowment, and n'= MRS, which decreases smoothly from b to l. Assuming true risk-neutrality in n, we have an induced utility function in g 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 utility 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).

5. Discussion

The concept of phlogiston, first suggested by Greek philosophers, entered the physics 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 was quite intuitive and able to organize many disparate physical phenomena such as combustion of charcoal (it released phlogiston) and smelting of metal ores (the metal absorbed phlogiston). Yet the phlogiston concept produced few novel predictions and the fluid was never isolated in the laboratory. The concept never became operational, and it faded after the emergence of Lavoisier's oxidation/reduction theory in the late 1780s (McKenzie, 1960, chapter 6).

Is non-linear expected utility the phlogiston of 20th century economists? First suggested by Gabriel Cramer and by Daniel Bernoulli (1738), expected utility theory entered the economic mainstream 60 years ago with the work of John von Neumann and Oscar Morgenstern (1944). The theory had intuitive appeal, spurred numerous theoretical elaborations, and seemed at first to organize many disparate economic phenomena such as insurance, financial asset prices, and gambling. But our reading of the empirical literature indicates that it has not yet fulfilled its early promise. As shown in section 2 above, there are no stable and consistent estimates of individual utility functions, despite numerous attempts in the last 50 years. Section 3 shows that the explanations of general economic phenomena based on non-linear utility have not aged well.

Phlogiston theory did not disappear when it encountered awkward facts, such as having to have negative mass). It vanished from respectable science only after a better theory came along.[17] Even if all our criticisms are acknowledged, expected utility theory will survive till economists are convinced that they have something better to replace it.

Behavioral economics might seem to be the most promising source for a better alternative. This popular new field takes the existence of non-linear utility functions as a given, and is preoccupied with finding specific non-linear utility functions that might fix the empirical inadequacies of expected utility theory. Most of the behavioral economics work is thus subject to our critique and is included in the material we discussed in section 2. Indeed, many of the recent studies cited above were inspired by behavioral economics. But we are skeptical about the theoretical alternatives proposed so far by behavioral economists.

The central model in behavioral economics seems to be Kahneman and Tversky's (1979) prospect theory. The prospect theory postulates an S-shaped value function u similar to a Markowitz utility function: it is convex below an inflection point z and concave above. It has at least 4 free parameters, including the inflection point (or reference point)[18] and slope and shape parameters for upper and lower ranges of outcomes.[19] Fitting these additional free parameters only compounds the inconsistency problems across individuals and tasks we highlighted earlier.

---

[17] 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."

[18] Possibly insert the story here of Kahneman (anonymized?) insisting on retaining it as a free parameter.

[19] 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., $a(x)>0$ for $x> z$ ["risk aversion for gains"] and $a(x)<0$ for $x< z$["risk seeking for losses"].

Despite the great flexibility derived from its free parameters, the value function seems inconsistent with some observed behavior. For example, purchasing insurance is difficult to reconcile with risk seeking for losses. Prospect theory then invokes a probability weighting function w; the DM is hypothesized to maximize $\sum w(p_i)u(x_i)$ rather than $\sum p_i\, u(x_i)$. The weighting function also has an inflection point separating concave and convex segments. When behavior is found inconsistent across tasks with the identical probabilities and outcomes, then prospect theory invokes an unmodelled phase of editing and adjustment. If prospect theory doesn't fit the data well, there are a host of alternatives such as regret theory, rank-dependent utility, etc. (Hey and Orme, 1994; Harless and Camerer, 1994). Clearly through its liberal use of free parameters, behavioral economics is capable of explaining almost anything ex post; its power to predict behavior and outcomes in novel situations will have to be demonstrated before it can be considered a serious candidate. **[unobservable free parameter discussion]**

Economists traditionally have distinguished themselves among social scientists by following an austere maxim: put the explanatory burden on potentially observable opportunities (e.g., prices and incomes), not on unobservables (e.g., arbitrary preferences or beliefs). The maxim has often led to distinctive predictions and new insights (e.g., Stigler, 1984). Behavioral economics, in pursuing an agenda based on unobservable explanatory variables, fails the observables and austerity test.

Our suggestion on the way forward is to use linear utility which makes EU and EV criteria identical, and look more carefully at the opportunity set, especially at how the possible outcomes of the decision interact with future opportunities and past commitments. For example:

- In analyzing the demand for insurance policies, don't assume a distribution of unobservable concavity parameters that produces the observed demand via maximization of expected utility. This would short circuit the analysis. Instead, consider the potentially observable patterns of commitments and embedded options. What additional costs do people incur when they suffer losses from fire, theft, etc.? For example, one might predict that homeowners with larger mortgages will carry more life insurance, and less discretionary fire insurance. (It's hard to imagine a straightforward EU approach for life insurance..)
- In analyzing gambling, don't assume a distribution of unobservable convexity parameters that produces the observed behavior via maximization of expected utility. Instead, consider potentially observable bailout options, etc. 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 him clout as well as wealth, while loosing big would only reinforce his current status without a serious threat to his survival.

Of course, not all new predictions will hold up in empirical tests. We think that behavioral economics (or cognitive psychology, or learning models) has the comparative advantage in explaining the behavior of inexperienced DMs in novel situations with low stakes. With greater incentives and learning opportunities, DM choices are more likely to be settled and coherent, and traditional economic analysis such as ours has its best shot.[20] Prediction failures in such settings are far more damaging to our proposed approach.

[20] Friedman (1989) develops this point explicitly, and shows how S-shaped value function for an inexperienced DM gradually converges to the true utility function (e.g., linear) as experience accumulates.

Our proposal involves estimates of net payoff as a possibly non-linear function of gross payoff. But what is net payoff? In principle it represents the DM's consumption opportunities, broadly conceived, and is usually proxied by expected terminal wealth. But this definition raises several conceptual issues with practical implications.

- The proxy ignores the fact that the more urgent consumption items (e.g., feeding one's family) will be met with the first dollars of terminal wealth and less urgent desires (e.g., vacation extras) with subsequent dollars. This is the classical argument for diminishing marginal utility of money, or concave utility. Arrow, and Friedman and Savage DATE??, brushed aside such arguments, but they strike us as reasonable.
- Hence we do not assert that "true" utility is precisely linear in terminal wealth. For the sake of parsimony, however, we do recommend treating it as if it were. We believe the linear approximation will 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 $1000 today implies a lifetime gain or loss of only a few cents in daily consumption.
- This perspective gives financial market frictions an important role. The Masson DATE??argument shows that if frictions dissuade a risk neutral DM from adjusting future consumption plans, then the net payoff is a concave (or perhaps more complex) function of the gross current period payoff.
- Another conceptual issue is whether net income (as in the examples of section 4) is a good proxy for terminal wealth. Assuming (approximate) linearity, we sidestep the debate over Rabin's (2000) calibration: does EU apply to terminal wealth, or income, or both? (Cox and Sadiraj, 2002) With linear (or CARA) utility, it doesn't matter.
- The main practical issue is whether the estimated net utility function fully captures the impact on terminal wealth. Portfolio effects of all sorts are relevant, although our emphasis has been on embedded options and preexisting obligations.
- Another perspective on the same point is that we assume the net income is separable from other components of terminal wealth. Appendix A shows that only linear utility is robust to deviations from this assumption.

As a scientific theory, non-linear utility theory has a remarkable track record. Despite 50 years of failed attempts to operationalize its basic construct, the non-linear utility function, the theory continues to dominate the economic analysis of risky choice. Economists have not abandoned the approach, but instead have proposed ever more complicated variants: utility functions with kinks, non-linear probabilities, transformation via the distribution function (or rank-dependence), and numerous other transformations and variants. In our view, such proposals have diverted attention from empirically more useful avenues of research.

We propose to return to the roots of choice theory, and put the explanatory burden on potentially observable opportunities rather than on arbitrary, unobservable utilities and beliefs. The foundations of finance are being reconstructed using options theory instead of risk aversion, and we believe parallel efforts hold great promise for the rest of economics.

Figure 0: Evidence from Stock Market on Risk and Return
(Source: Prepared by authors from data in Black (1992, Exhibits 3 and 4)

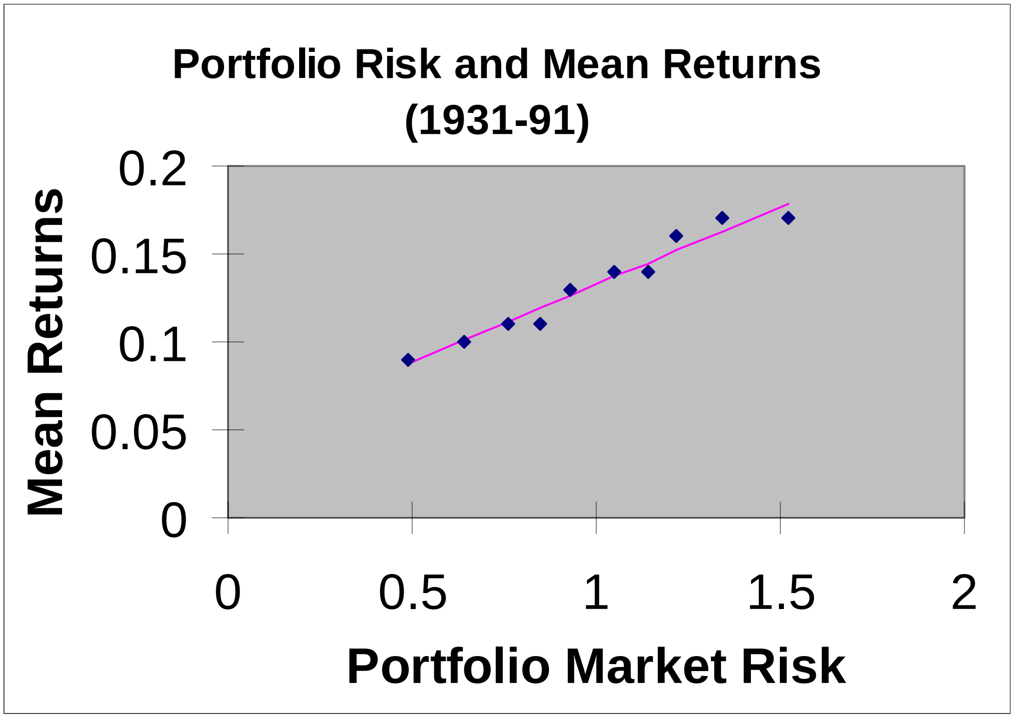

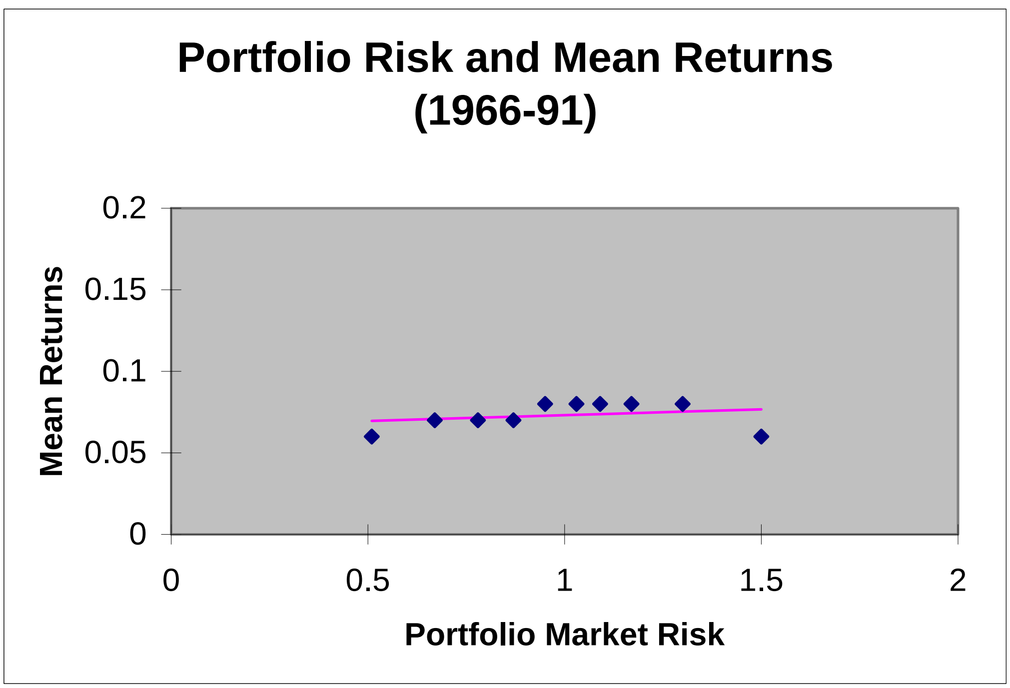

Figure 1: Concave Net Payoff

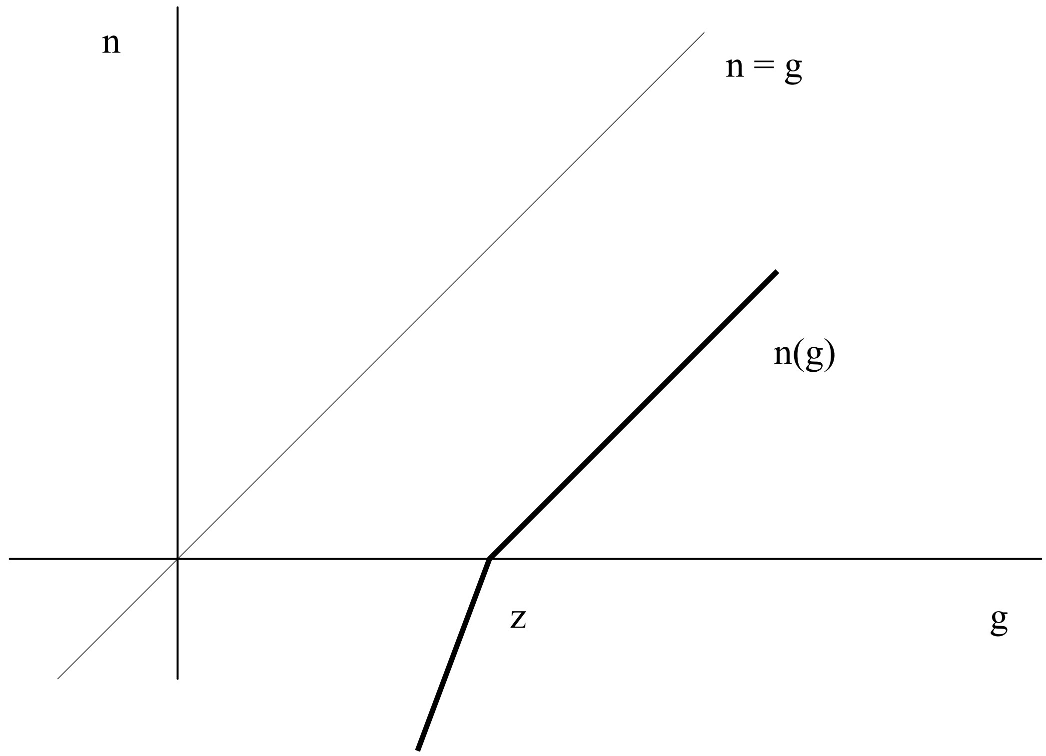

Figure 2: Tournament Payoff

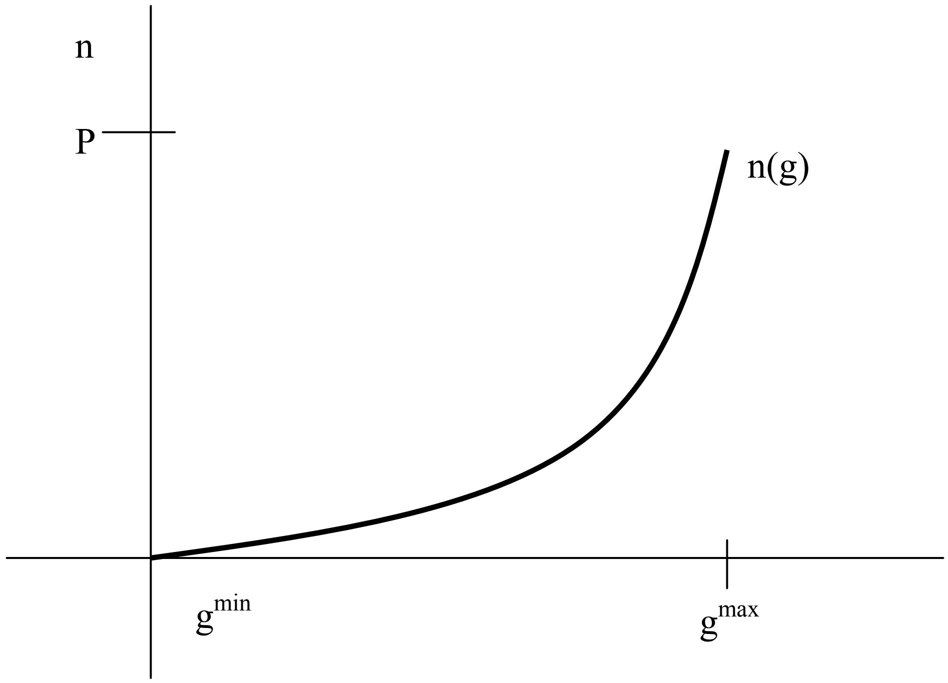

Figure 3: Convex Net Payoff

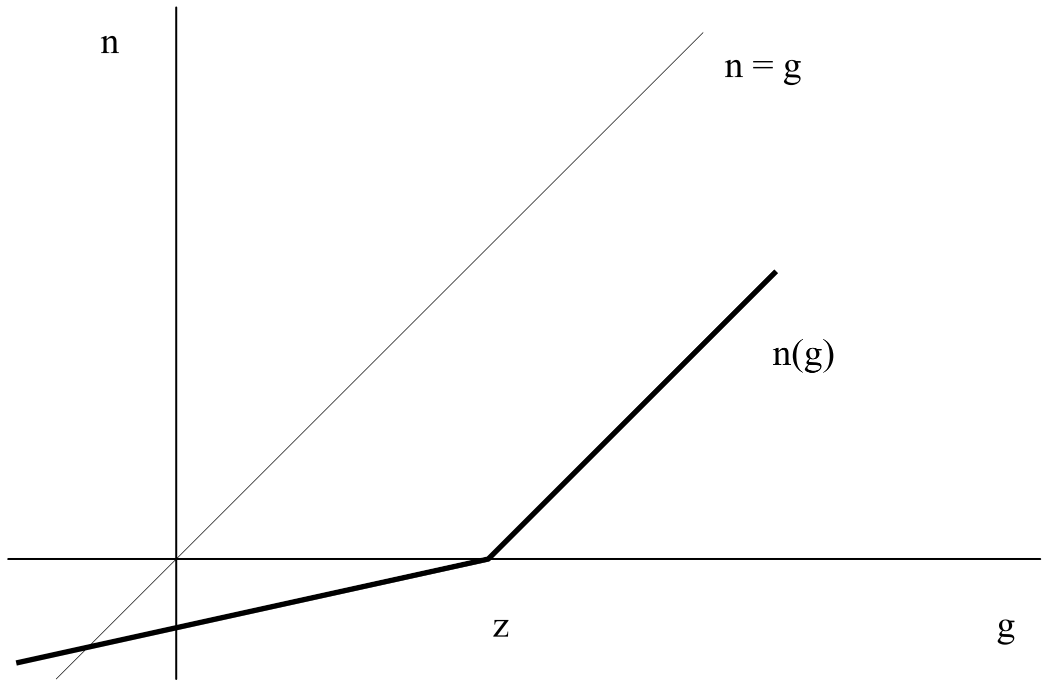

Figure 4: Payoff with Means-Tested Subsidy

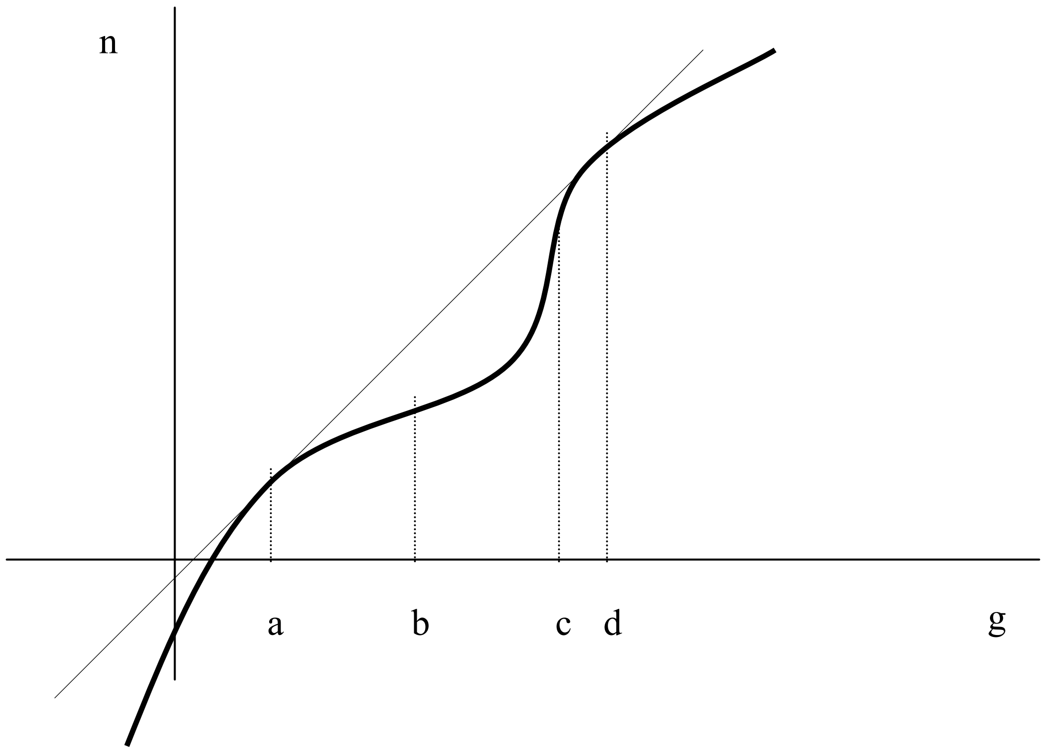

References

Ait-Sahalia, Y., and A. W. Lo. 2000. "Nonparametric Risk Management and Implied Risk Aversion," Journal of Econometrics 94, 9-51.

Altman, Edward I. 1989. "Measuring Corporate Bond Mortality and Performance," The Journal of Finance 44:4 (September), 909-922.

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.

Arrow, Kenneth J. 1971. Essays in the Theory of Risk Bearing. Amsterdam: North Holland.

Bank of International Settlements. 1999. Principles of the Management of Credit Risk. Basel Committee Publications No. 54. http://www.bis.org/publ/bcbs54.htm.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," Quarterly Journal of Economics, CXII (1997), 537-580.

Becker, Gary, Kevin Murphy, and Iván Werning, "Status, Lotteries and Inequality," University of Chicago mimeo (2000).

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.

Black, Fischer. "Beta and Return," Journal of Portfolio Management 20:8-18 (Fall 1993).

Borch, K. 1962. "Equilibrium in a Reinsurance Market," Econometrica pp. 424-444.

Brav, Alon, George Constantinides, and Christopher C. Geczy. 2002. "Asset Prices with Heterogenous Consumers and Limited Participation: Empirical Evidence," Journal of Political Economy ??.

Brunnermeier, Markus, "On Bounded Rationality and Risk Aversion," LSE Financial Markets Group Discussion Paper 255 (1997).

Binswanger, Hans P. 1978. "Attitudes toward Risk: Implications for Econimic and Psychological Theories of Behavior under Risk of an Experiment in Rural India," New Haven, CT: Yale Economic Growth Center Disc. Paper 286.

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.

Carney, R.E. 1971. Risk-Taking Behavior: Concepts, Methods and Applications to Smoking and Drug Abuse. Springfield, IL: Charles T. Thomas.

Chan, Louis K., and Josef Lokanishok, "Are Reports of Beta's Death Premature?" Journal of Portfolio Management, (Summer 1993), pp. 51-61.

Chipman, John S., Leonid Hurwicz, Marcel K. Richter, and Hugo F. Sonnenschein, eds. 1971. Preferences, Utility, and Demand. New York: Harcourt Brace Jovanocich.

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)

Chetty, Raj. 2003. "A New Method of Estimating Risk Aversion." University of California, Berkeley and NBER Working Paper (http://elsa.berkeley.edu/~chetty/papers/curvature.pdf)

Cordell, David M. 2001. "RiskPACK: How to Evaluate Risk Tolerance," Journal of Financial Planning (June).

Dickerson, Mark G. Compulsive Gamblers (New York: Longman, 1984), p. 40.

Dreze, Jacques, Essays on Economic Decisions Under Uncertainty (Cambridge: Cambridge Univ. Press, 1987).

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, "Probability Preferences in Gambling," American Journal of Psychology 66, 349-364.

Edwards, Ward, "The Prediction of Decisions among Bets," Journal of Experimental Psychology 50:3 201-214.

Ezell, John Samuel. Fortune's Merry Wheel: The Lottery in America (Harvard University Press, Cambridge, 1960), p. 1.

Fair, Ray C. 2003. "Risk Aversion and Stock Prices," Cowles Foundation and International Center for Finance, Yale University, Mimeo.

Fama, Eugene and Kenneth French, "The Cross Section of Expected Stock Returns," Journal of Finance 47:427-465 (June, 1992).

Fisher, Lawrence. 1959. "Determinants of Risk Premiums on Corporate Bonds." Journal of Political Economy 67, 217-237.

Friedman, Milton. 1953. Positive Theory essay.

Friedman, Milton, and Leonard J. Savage. 1948. "The Utility Analysis of Choices Involving Risk," Journal of Political Economy 56 (August), 279-304.

Friedman, Milton, and Leonard J. Savage. 1948. "The Utility Analysis of Choices Involving Risk," Journal of Political Economy 56 (August), 279-304.

Gertner, Robert, "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'," Quarterly Journal of Economics, 108 (1993), 507-22.

Grossman, S. J. and Robert Shiller. 1982. "Consumption correlatedness and risk measurement in economies with non-traded assets and heterogenous information," Journal of Financial Economics 51, 1379-1403.

Guiso, L. and Paiella, M. 2001. 'Risk Aversion, Wealth and Background Risk'. CEPR Discussion Paper no. 2728. London, Centre for Economic Policy Research. http://www.cepr.org/pubs/dps/DP2728.asp.

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.

Harlow, W. and K. Brown. "Understanding and Assessing Financial Risk Tolerance: A Biological Perspective." Financial Analysts Journal. November/ December 1990: 50–80.

Harris, Robert S. and Felici C. Marston, "Value versus Growth Stocks: Book to Market, Growth and Beta," Financial Analysts Journal (September-October, 1994).

Harrsion, 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.
Harrsion, Glenn W., Morten Igel Lau, E. Elisabet Rutström and Melonie B. Williams, “Eliciting Risk and Time Preferences Using Field Experimentas: Some Methodological Issues,” University of Central Florida Working Paper, August 2003.
Hartley, Roger and Lisa Farrell, “Can Expected Utility Theory Explain Gambling?” American Economic Review 92 (2002), 613-624.
Haugen, Robert A., The New Finance: The Case Against Efficient Markets (Englewood Clifffs, N.J.: Prentice Hall, 1995)
Henrich, Joseph and Richard McElreath, 2002. "Are Peasants Risk-Averse Decision Makers?" Current Anthropology 43:1 (February), 172-181.
Hey, John D., “Experimental Economics and the Theory of Decision Making Under ncertainty,” Geneva Papers on Risk and Insurance Theory, 27(1), June 2002, 5-21.
Hickman, W. Braddock. 1958. Corporate Bond Quality and Investor Experience. Princeton, N.J.: Princeton University Press.
Holt, Charles A., and Laury, Susan K., “Risk Aversion and Incentive Effects,” American Economic Review, 92(5), December 2002, 1644-1655.
Ibbotson Associates, Encore Database, 1998.
Issac, Mark, and James, .2000. « , » Journal of Risk and Uncertainty
Jackson, D. N., L. Hourany, and N. J. Vidmar. 1972. « A Four Dimensional Interpretation of Risk Taking, »Journal of Personality 40, 483-501.
Jackwerth, Jens Carsten. 2000. « Recovering Risk Aversion from Option Prices and Realized Returns, » The Review of Financial Studies 13 :2, 433-451.
Jeffery, R. W. 1989. “Risk behaviors and health: contrasting individual and population perspectives,” American Psychologist 44, 1194-1202.
Jianakoplos, Nancy Ammon, and Bernasek, Alexandra, “Are Women More Risk Averse?” Economic Inquiry, 36, October 1998, 620-630.
Kahneman, Daniel, Peter P. Wakker, and Rakesh Sarin, “Back to Bentham? Explorations of Experienced Utility,” Quarterly Journal of Economics, CXII (1997), 375-405.
Kandel, Shmuel and Robert F. Stambaugh, “Asset Returns and Intertemporal Preferences,” Journal of Monetary Econmics, 27 (1991), 39-71.
Knight, Frank H. 1921. Risk, Uncertainty and Profit. Boston: Houghton Mifflin & Co.
Kocherlakota, Narayana, “The Equity Premium: It’s Still a Puzzle,” Journal of Economic Literature, 24 (1996), 42-71.
Kothari, S.P., Jay Shanken, and Richard G. Sloan, “Another Look at the Cross-Section of Expected Stock Returns,” University of Rochester Working Paper, 1992.
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.
Leigh, Barbara C. 1999. “Peril, Chance, Adventure: Concepts of Risk, Alcohol Use and Risky Behavior in Young Adults.” Addiction 94:3, 371-383.
Leland, J.W. and J. Grafman. 2003. “Experimental Tests of the Somatic Marker Hypothesis,” National Institute of Neurological Disorders and Stroke Working Paper.
Levin, Irwin P.; Snyder, Mary A., and Chapman, Daniel P., ‘‘The Interaction of Experiential and Situational Factors and Gender in a Simulated Risky Decision-Making Task,’’ Journal of Psychology, 122(2), March 1988,173–181.

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.
Lipsitt, L. P., and L. L. Mitnick, eds. 1991. Self-Regulatory Behavior and Risk-Taking: Causes and Consequences. Norwood, N.J.: Ablex.
List, John A. 2004. “Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace.” Econometrica 72:2 (March), 615-625.
Machina, Mark, “Choice Under Uncertainty: Problems Solved and Unsolved,” Journal of Economic Perspectives, 1 (1987), 121-154.
Markowitz, Harry, “The Utility of Wealth,” Journal of Political Economy, 60, April 1952, 151-158.
Markowitz, Harry. 1952. “Portfolio Selection,” Journal of Finance 7 (March), 77-91.
Marshall, John M. “Gambles and the Shadow Price of Death,” American Economic Review 74:1, 73-86.
MacCrimmon, K. R. and D.A. Wehrung. 1986). Taking Risks: the management of uncertainty. New York: The Free 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.
Masson, Robert T. 1973. “Utility Functions with Jump Discontinuities: Some Evidence and Implications from Peasant Agriculture,” Western Economic Journal XII:4 (December), 559-566.
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.
Metrick, Andrew, “A Natural Experiment in Jeopardy!,” American Economic Review, 85 (1995), 240-253.
Miller, Merton H. 1999. “The History of Finance.” The Journal of Portfolio Management 25:4 (Summer), 95-101.
Mosteller, F. and P. Nogee. 1951. “An Experimental Measurement of Utility,” Journal of Political Economy 59 (October), 371-404.
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.
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 (1981)
Robson, Arthur J., “The Evolution of Attitudes to Risk: Lottery Tickets and Relative Wealth,” Games and Economic Behavior, 14, 1996, 190-207.
Roll, Richard. :A Critique of the Asset Pricing Theory’s Tests; Part 1: On Past and Potential Testability of the Theory,” Journal of Financial Economics 4:129-176 (March 1977).
John Rose Crance, Gambling Without Guilt: The Legitimatisation of an American Pastime. (Pacific Grove: Brooks/Cole Publishing Co., 1988), p. 54.
Rosenberg, Barr, Kenneth Reid and Ronald Lanstein, "Persuasive Evidence of Market Inefficiency," Journal of Portfolio Management (11, 1985).
Royden, H. L. Real Analysis, Second Edition. New York:MacMillan, 1968.
Rubin, P. H., and Paul, C. W., “An Evolutionary Model of Taste for Risk,” Economic Inquiry, 17, 1979, 585–595.
Savage, Leonard J. 1954. Foundations of Statistics. New York: John Wiley & Son.
Schilbred, Cornelius M. 1973. “The Market Price of Risk,” The Review of Economic Studies 40:2 (April), pp. 283-292.
Schmidt-Nielsen, Knut. 1997. Animal Physiology: Adaptation and Environment. Cambridge University Press.
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.
Sinn, Hans-Werner. 2002. “Weber’s Law and Biological Evolution of Risk Preferences: The Selective Dominance of the Logarithmic Utility Function.” CESifo Working Paper #770.
Smart, Scott B., William L. Megginson and Lawrence J.Gitman 2004. Corporate Finance. Cincinnati, OH: Sotu-Western.
Stattman, Dennis, "Book Values and Stock Returns," The Chicago MBA (4, 1980)
Szpiro, George, “Measuring Risk Aversion: An Alternative Approach,” The Review of Economics and Statistics, 68 (1986), 156-159.
TIAA-CREF. 1997. Building Your Portfolio with TIAA-CREF. New York: TIAA-CREF.
Ukhov, Andrey D. 2003. “Time Varying Risk Aversion: Evidence from Russian Lottery Bonds,” in Three Essays in Financial Economics, Doctoral Dissertation, Yale University.
Yaari, Menachem. 1987. “The Dual Theory of Choice Under Risk,” Econometrica 55:1, (January) 95-115.
Yook, Ken C. and Robert Everett. 2003. Journal of Financial Planning (August)
Wagenaar, W., 1988. Paradoxes of Gambling Behaviour. London: Erlbaum.
Wallace, A. “Is Beta Dead?” Institutional Investor, 14:22-30 (July 1980).
Wolf, Charles and Larry Pohlman, “The Recovery of Risk Preferences from Actual Choices,” Econometrica, 51 (1983), 843-850.
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.