The Psychology of Human Risk Preferences and Vulnerability to Scare-Mongers: Experimental Economic Tools for Hypothesis Formulation and Testing

Glenn W. Harrison[1]

Don Ross[2 *]

October 2016

1 – Department of Risk Management and Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA. Harrison is also affiliated with the School of Economics, University of Cape Town, South Africa.

2 – Department of Philosophy, University College Cork, Ireland; School of Economics, University of Cape Town, South Africa; and Center for Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA.
* Corresponding author, email: don.ross931@gmail.com

Abstract

The Internet and social media have opened niches for political exploitation of human dispositions to hyper-alarmed states that amplify perceived threats relative to their objective probabilities of occurrence. Researchers should aim to observe the dynamic "ramping up" of security threat mechanisms under controlled experimental conditions. Such research necessarily begins from a clear model of standard baseline states, and should involve adding treatments to established experimental protocols developed by experimental economists. We review these protocols, which allow for joint estimation of risk preferences and subjective beliefs about probabilities and their distributions. Results we have obtained on such estimates, from populations in various countries, are gathered for comparison. Most people show moderate risk aversion in non-alarmed states. We also find universal heterogeneity in risk preference structures, with substantial sub-samples weighting probabilities in such a way

as to display “probability pessimism” (rank dependent utility), while others make risky choices in accordance with expected utility theory.

Keywords

security threats, experimental economics, human risk preferences, subjective beliefs about probabilities, heterogeneity of preference and belief structures, expected utility theory, rank dependent utility

1. Introduction

A currently salient focus of challenges to governance models stems from opportunistic political exploitation of evolved human threat-response systems. Woody and Szechtman [2016, this volume] argue persuasively that new social media facilitate this exploitation, and lead to widespread perception of risk that bears little relationship to objective dangers, particularly to dangers that lie within the plausible control of the hyper-alarmed individuals. In this context it is useful to review what is known about the relationship between objective risks and subjective risk attitudes in people’s *non*-alarmed, or “resting,” states. An obvious avenue for future experimental research is to observe the dynamic “ramping up” of security threat mechanisms under controlled conditions. Such research necessarily begins from a clear model of standard baseline states. Furthermore, such work will be most informative to the extent that it adds new treatments to the experimental protocols that have been used for probing the baseline relationships between subjective and objective risk. The primary disciplinary site of this protocol development has been experimental economics.[1]

Most decisions made by humans and other animals involve risk or uncertainty. In a typical situation of choice, which available alternative response is optimal for a decision maker (DM), in terms of either organismal fitness or individual utility, is conditional on which member of a set of possible states of the world obtains, where the probability of each state in the set is neither 0 nor 1. Choices by DMs under such circumstances can be interpreted as expressing either virtually or literally computed *subjective estimates* of the relevant probabilities, multiplied by the relative costs and benefits to the DM associated with each possible decision outcome. Following standard terminology, a choice is said to involve *risk* when the information that is required to accurately estimate relevant objective probabilities, given the computational resources available to the DM, is available in principle, regardless of whether it is actually harvested or used by the DM. If some such information is not available then the choice is said to involve *uncertainty*.[2]

Consider an example. A subordinate male social animal faces a choice between attempting to mate, or not, with a specific female on a particular occasion. Suppose that if this attempt is observed by a dominant male, punishment will result. For simplicity, suppose that observed mating attempts by subordinates are always punished, i.e., that the probability of punishment given observation of a mating attempt is 1. If probability of observation varies with circumstances, then whether it is optimal for the subordinate to make a specific attempt depends on multiplying this probability by the *expected payoffs* to the subordinate

[1] For a survey of pre-experimental background theory in economics, see Gollier [2001].

[2] A further distinction is often drawn between situations in which objective probabilities could be calculated given more information, and situations in which alternative states are insufficiently specified for their probabilities to be calculated even given full information. Choices in conditions of the latter type are said to involve *ambiguity*. We will set this distinction aside for present purposes.

associated with, respectively, successful and unsuccessful mating attempts. The payoffs in question may be positive or negative with respect to any particular *ex ante* reference point. At the most general level of scientific abstraction, payoffs can be represented as changes in relative expected fitness coefficients. In the case of an organism that can pursue idiosyncratic interests that diverge from expected fitness, costs and benefits may instead be measured in terms of ordinal or cardinal *expected utility* (von Neumann and Morgenstern [1944]; Savage [1954]). In the example, expected utility from successful mating, from spurned mating attempts, and from hostile response by a dominant observer would incorporate material, social, and psychological rewards and punishments. Choice options characterized in terms of expected payoffs are referred to as *prospects*.

Humans had some ancestors that lacked the cognitive resources to explicitly compute probabilities associated with risky and uncertain choices. That the genes of these ancestral lines persisted is sufficient evidence for the proposition that they had evolved dispositions to respond differentially to varying objective probabilities.[3] These dispositions would have accurately optimized expected payoffs only in conditions that occurred with sufficient past frequency to have been tracked by adaptive selection, and that were distinguishable by the ancestors' perceptual sensitivities. Since this implies common gaps between ideal and actual conditions, cognitively unsophisticated human ancestral DMs sometimes made choices that involved error, in the sense that a stochastically dominated prospect[4] was chosen. Following standard reasoning in evolutionary psychology, we assume that traces of these evolved dispositions remain present in cognitively normal human adults

[3] Some of these would have been dispositions for conditioned learning of recurrent contingencies.
[4] An option of type O stochastically dominates an option of type O' if and only if the statistically expected payoff of an instance of O is greater than the statistically expected payoff of an instance of O'. This allows for occasions when the *realized* payoff from a specific instance of O' is higher than that of the alternative O-type option.

and children, where they manifest as *biases*, that is, as tendencies to make certain kinds of errors more than other kinds of errors.

This standard reasoning is silent on an important, general, empirical question. Did natural selection produce a standard set of biases common in all cognitively normal people, or a polymorphism of biases? Formulating possible answers to this question requires recognition of the complex *structure* of risky and uncertain choices. Biases might be identified at any or all of several levels of choice problem specification, of which three have received special emphasis in economics:

(1) *Risk preferences*: Does a person have a*certainty equivalent* (CE) for a prospect that is greater than, equal to, or less than the expected payoff of the prospect? A CE is the non-risky payoff, or "sure thing," that is equivalent, for the DM, to the risky prospect. A DM with a CE that is less (greater) than the expected payoff is said to be *risk averse* (*loving*) with respect to that prospect context in question. aA DM is *risk neutral* in the prospect context if she has a CE that is equivalent to the risky expected payoff.

(2) *Subjective belief priors*: Does a person subjectively represent probability distributions of risky prospects in ways that systematically depart from, but co-vary with, objective probability distributions? A DM that tends to either over-weight or under-weight probabilities depending on where they lie in objective probability distributions is said to *probability weight,* and is characterized as having *rank-dependent* attitudes to risk in the relevant prospect context (Quiggin [1982]).

(3) *Reference point sensitivity*: Do people try to avoid what they perceive as "backward[5]" changes in utility wealth *per se*? A DM that varies risk preferences and/or subjective belief priors systematically with respect to a *reference point* that separates prospects into classes subjectively framed as "gains" or "losses," where prospects falling on the "loss" side of the reference point are weighted more heavily than prospects on the "gain" side of the reference point, is said to be *loss averse*.

Experimental economists gather data on these and other dimensions of variation in choice under uncertainty among people (and other animals; see Kagel, Battalio and Green [2007]) by presenting subjects with choices among prospects, typically but not exclusively as "lotteries" over monetary rewards, that are designed in such a way is to elicit patterns in data that can be estimated using a range of formal theories. *Expected utility theory* (EUT) (Savage [1954]; Binmore [2009]) describes choices that would be made by a probability un-biased prospect optimizer. Choices made by DMs with rank-dependent, probability weighted preferences can be characterized by a family of models descended from Quiggin [1982] that are collectively referred to as *rank dependent utility theory* (RDU). Choices reflecting loss aversion dependent on a subjective reference point, which also incorporate rank-dependence, are currently modeled using *cumulative prospect theory* (CPT), due to Tversky and Kahneman [1992]. It is unfortunately common to find behavioral scientists who are not economists, along with some careless economists, assuming that people in general manifest the biases identified by CPT. It casts no aspersion on the importance of CPT as one item in the

---

[5] This refers to a folk intuition, not directly built into the axioms of any formal choice theory, that people try to avoid situations in which they have less wealth, including less utility wealth, tomorrow than they have today.

theorist's toolbox to point out that this assumption is empirically unsupported (Harrison and Swarthout [2016]).

EUT, RDU, and CPT structurally identify *processes,* or classes of processes, that generate *observable choice data*. The processes in question are sometimes literally computed, either by algorithms implemented in neural circuitry or by people using external devices that record and compile operational steps. But more often they are "virtual" in the sense associated with Dennett's [1987] intentional stance. That is, a formally characterized decision making process typically denotes an equivalence class of relationships between prospect contexts and observed choices in which operations distinguished in the theorist's analysis do not necessarily represent discrete operations performed by the DM, but might rather represent elements of the economic, social, or biological problem as characterized by an analyst modeling its solution. To the extent that parallel distributed processing or "connectionist" models of decision making approximate the dynamics of neural processing, the virtual-process interpretation may be the standard default one (Clark [1989]) for decisions processed by a "raw brain" functioning on its own. On the other hand, humans make many of their risky decisions in social contexts where they use one another as sources of imitation and consultation, and where they furthermore have recourse to many kinds of institutional "scaffolding" such as actuarial tables, investment guides, books, the internet, etc. (Hutchins [1995], Clark [1997]). In such cases literal-process interpretations are more plausible, though still require special direct evidence regarding implementation mechanisms.

Recent experimental economic research has established an important, general fact for evolutionary psychologists to investigate and model. This is that human risky decision making involves *structurally specifiable, predictable heterogeneity across individuals, across choice domains,* *and across decision environments*. That is, it is *neither* characterized by uniform biases that affect all people in all choice contexts, *nor* is it chaotic and unpredictable. This has immediate implications for policy design and governance. A policy or set of institutionally constructed incentives that guides some part of a randomly selected human population toward one risky choice tendency will produce predictably divergent effects in other parts of such a population. Economists in their role as policy advisors are directly engaged with these effects. As part of the project of studying hyper-alarmed security threat systems, evolutionary psychologists should also study the biological and cultural origins of the multi-faceted heterogeneity discovered by economists, which may be expected to shed light on the extent of the system's adaptability and plasticity.

Section 2 reviews the dimensions along which heterogeneity in human risk responses has been revealed through economic experimentation. Section 3 presents the family of basic theories that are used to represent a specific, but fundamental, aspect of this heterogeneity, namely variation with respect to the structure of risk preferences. Observations of the relative frequencies of these varying structural types, derived from our experimental work, are then summarized in tabular form. This is the first time this suite of studies have been presented together so as to reveal recurrent patterns across different study populations. In Section 4 we turn to the theory and evidence behind a second source of heterogeneity, variation in subjective beliefs about probabilities. Section 4 concludes with some reflections on the significance of the reported observations for anticipated future research by psychologists and others, including economists who can supply complementary formal theory and econometric estimation tools, on the vulnerabilities that, under some circumstances, render people prone to being roused into states of hyper-alarm about abstract threats that is out of proportion to objective risk.

2. Dimensions of Heterogeneity

Economists are interested in heterogeneous response to risk and uncertainty for more than just general theoretical and descriptive purposes. They also must contend with it for practical purposes. For example, should women in a specific sample be predicted to display more risk aversion than men in that sample? Policy alternatives on which economists are asked to offer advice are often tested by being administered as randomized treatments in experiments. However, in light of heterogeneity, randomization to treatment is often insufficient to ensure the absence of confounding differences in samples. Heterogeneity of subject response directs attention to causal mechanisms generating it, which might interact with the treatment. This in turn implies attention to structural features of models, and complicates the statistical issues around randomized evaluation; see Heckman [1992], Keane [2010a], [2010b], and Leamer [2010]. The very point of worrying about sample selection is recognition of the possibility of effects of heterogeneity, so an experimenter cannot be casual about controlling for it. The matter is particularly acute when selection is on risk attitudes of treated and untreated subjects, so-called "randomization bias" (Harrison, Lau and Rutström [2009]; Harrison and Lau [2014]).

Heterogeneity is subtle. It can refer to differences across individuals at a point in time, conditional on a model of decision-making. It can refer to differences over time for the same individual and model of decision-making. It can refer to differences in the model of decision-making for an individual. And it can refer to differences in the behavior of the same individual across domains (e.g., choices over financial risks and choices over health risks).

A. Heterogeneity Across Individuals for a Given Model

There are three broad approaches to allowing for individual heterogeneity in experimental choice data.

The first is to collect enough observations from each subject to be able to directly estimate models for that individual. Hey and Orme [1994] remains a leading exemplar of this approach, which has become more popular in recent applications. A main practical limitation to this method is that as one adds experimental treatments to test hypothesized special effects, power may quickly be compromised given limited experimental resources.

The second approach is to collect a standard list of observable demographic characteristics from each individual, pool the responses across all subjects, and condition estimation of parameters on the observed characteristics. Thus one might estimate a model in which a risk aversion parameter is a linear function of sex, age, race, income, and so forth. In this case one obtains estimates of the coefficients of the effect of each of the characteristic on the parameter, as well as a constant term (e.g., Harrison, Lau and Rutström [2007]). The danger with this approach is that there may be some characteristic of individuals that is not observed by the experimenter and that systematically affects responses: this is the so-called "unobserved individual heterogeneity" problem.[6]

The third approach to allowing for individual heterogeneity is to use econometric methods that are relatively agnostic about the form that heterogeneity takes, but that explicitly allow for it. This approach might be usefully termed the "random coefficients" approach. Instead of assuming that a risk aversion coefficient is a linear function of observables, one might assume that the coefficient is normally distributed across the sample.

[6] We say so-called, since the term "unobserved" is not literally correct. We do observe that a certain set of observations were generated by one subject, and another set by other subjects, and this observation is used in all standard methods to correct for unobserved individual heterogeneity. A related issue is the interaction of models of "behavioral error" and unobserved heterogeneity, identified in Ballinger and Wilcox [1997].

In effect, each subject is assumed to have some true coefficient value, but these values are viewed by the analyst as being distributed across the sample in a way that can be characterized by a normal distribution. One then estimates *hyper-parameters* to reflect this population distribution.[7]

For example, if one assumes a normal distribution, one would estimate a mean (population) risk aversion coefficient and a standard deviation in the (population) risk aversion coefficient. Each hyper-parameter is estimated by a point estimate and a standard error.[8] As the sample size gets larger, one would expect the estimated standard *errors* of the *hyper-parameters* to shrink, but one would not necessarily expect the standard *deviation* of the population *parameter* to shrink. This approach generalizes naturally to non-normal distributions for the population parameter, and to multivariate distributions where there are more than one population parameter.

There is no necessary tension between the second and third approaches.[9] However, in practice one tends to see data evaluated using one or the other method.

B. Heterogeneity Across Time for the Same Individuals and Model

A closely related issue is the temporal stability, or temporal homogeneity, of risk preferences even when one uses the same elicitation procedure. It is possible to define

---

[7] The terminology across the econometric literature is not standard, so the expression "hyper-parameter" can have other meanings.

[8] To be pedantic, there are then four estimates: the point estimate and standard error of the mean of the population parameter, and the point estimate and standard error of the standard deviation of the population parameter.

[9] Nor is there tension between the first approach and the latter two. It is likely that one cannot obtain reliable estimates for some models for each individual across a sample, due to errant behavior by those individuals, but that the subsequent analysis requires some estimate. In this case one could complement the individual-level estimate and standard error with a predicted point estimate and standard error from one of the latter two approaches. Providing the nature of this stochastic imputation is properly accounted for, the "complete case analysis" can then proceed over the entire sample. This problem is likely to become practically significant as small-scale experiments are increasingly used to augment the data collected in large-scale surveys in the field.

temporal stability of preferences in several different ways, reflecting alternative conceptual definitions and operational measures. Each definition has some validity for different inferential purposes.

Temporal stability of risk preferences can mean that subjects exhibit the same risk attitudes over time, or that their risk attitudes are a stable function of states of nature and opportunities that change over time. It is quite possible for risk preferences to be stable in either, both or neither of these senses, depending on the view one adopts regarding the role preference stability takes in the theory. The temporal stability of risk preferences is one component of a broader set of issues that relate to the state-dependent approach to utility analysis. This is a perfectly general approach, where the state of nature could be something as mundane as the weather or as fundamental as the individual's mortality risk. Relevant states could also include the opportunities facing the individual, such as market prices and employment opportunities. Crucial to the approach, however, is the fact that all state realizations must be exogenous, or the model will not be identified and inferences about stability will be vacuous. For example, whether or not someone has chosen to be a mother would not typically be thought of as an exogenous state realization.

Problems arise, however, when one has to apply this approach empirically. Where does one draw the line in terms of the abstract "states of nature"? Many alleged violations of EUT amount to claims that a person behaved as if they had one risk preference for one prospect context and another risk preference for a different prospect context. Implicit in the claim that these are violations of EUT is the presumption that the differences across the two prospect contexts were not contingent on some state of nature over which preferences could be different. Similarly, should we deem the preferences elicited with an open-ended auction procedure to be different from those elicited with a binary choice procedure, such as in the much-cited preference reversals of Grether and Plott [1979], because of some violation of EUT or just some change in the state of nature? Of course, it is a slippery inferential slope that allows "free parameters" to explain any empirical puzzle by shifting preference attributions. Such efforts have to be guided by direct evidence from external sources, lest they become open-ended specification searches.[10]

These issues can be resolved by evaluation of longitudinal experimental data, where one is able to collect information on changes in observable states of nature over time. Examples of such studies include Andersen, Harrison, Lau and Rutström [2008] and Harrison and Lau [2014].

C. Heterogeneity Across Models for the Same Individuals

Since different models of choice behavior imply different characterizations of risk attitudes, it is important that we make some determination about which of these models is to be adopted. One of the enduring contributions of experimental economics is that we now have a rich set of competing models of behavior in many settings, with EUT, RDU and CPT as the leading tools for representing choices under risk. Debates over the validity of these models have often been framed as "horse races," with one theory being declared the winner on the basis of some statistical test in which the theory is represented as a latent process explaining the observed choice data. On this methodology, if one theory explains more of the data than another theory, it is adopted for application to new samples and the "loser" is discarded. The problem with this approach is that it does not recognize the possibility that several data-generating computational or other processes may coexist in a population.

[10] Stigler and Becker [1977; p.76] note the nature of the impasse: "an explanation of economic phenomena that reaches a difference in tastes between people or times is the terminus of the argument: the problem is abandoned *at this point* to whoever studies and explains tastes (psychologists? anthropologists? phrenologists? sociobiologists?)."

Recognizing that possibility has direct implications for the characterization of the structure of human (and other animal) risk attitudes.

Ignoring the possibility of polymorphic risk preference structures can lead to erroneous conclusions about the domain of applicability of each theory, and is likely an important reason why the "horse races" pick different winners in different domains. For example, EUT is sometimes thought to govern choices involving financial risk, while failing to accurately characterize choices involving viscerally tempting sensations (Ainslie [1992]). For purely statistical reasons, if we conjecture that there are two or more latent population processes generating an observed sample, we can make more appropriate inferences if the data are not forced to fit a specification that assumes one data-generating process operating across the whole sample.

Heterogeneity in responses is well recognized as causing statistical problems in experimental and non-experimental data. Nevertheless, allowing for heterogeneity in responses through standard econometric methods, such as fixed or random effects, is not helpful when we want to identify which people behave according to which theory, and when. Heterogeneity can be partially recognized by collecting information on observable characteristics and controlling for them in the statistical analysis, as noted above. But this approach only recognizes *heterogeneity within a given theory*. This may be important for valid inferences about the ability of the theory to explain the data, but it does not allow for *heterogeneous theories* to co-exist in the same sample.

An approach to heterogeneity that reflects a more empiricist attitude, as proposed and illustrated by Harrison and Rutström [2005], is to construct a "wedding" of the theories by specifying and estimate a grand likelihood function that allows each theory to co-exist and have different weights. Such a likelihood function is referred to as a *mixture model*. The data can then identify the relative degree of support in the data for each theory. The wedding is consummated by the maximum likelihood estimates converging on probabilities that apportion non-trivial weights to each theory.

Harrison's and Rutström's [2005] results are striking where risk preference theories are concerned: EUT and the direct ancestor theory of CPT (Kahneman and Tversky 1979])[11] share the stage, in the sense that each accounts for roughly 50% of the observed choices. Thus, to the extent that EUT and CPT imply different things about how one measures risk aversion, and the role of the utility function as against other constructs, assuming that the data are generated by one or the other theory can lead to erroneous conclusions. The fact that the mixture probability is estimated with some precision, and that one can reject the null hypothesis that it is either 0 or 1, also indicates that one cannot claim that the equal weight to these models is due to chance.

The main methodological lesson from this exercise is that one should not rush to declare one or other model as a winner in all settings.[12] One would expect that the weight attached to EUT would vary across task domains, just as it can be shown to vary across observable socio-economic characteristics of individual DMs. In fact, one can represent the conjecture that a mixture model could also characterize the same individual, either synchronically or diachronically or both, as in the "dual criteria" models of the psychologist Lopes [1984] (e.g., see Andersen, Harrison, Lau and Rutström [2014]). Alternatively, mixture models might characterize the same individual in different contexts: when there are lotteries defined over 1, 2 or 3 prospects, with rounded probabilities, versus lotteries defined over 4

[11] The findings extend to the CPT model of Tversky and Kahneman [1992].

[12] A concrete implication, considered at length in Harrison and Rutström [2005; §5], is that the rush to use non-nested hypothesis tests is misplaced. If one reads the earlier literature on those tests it is immediately clear that they were viewed as poor, second-best alternatives to writing out a finite mixture model and estimating the weights that the data place on each latent process. The computational constraints that made them second-best decades ago no longer apply.

or more prizes with probabilities presented to the 3rd decimal place (e.g., Wilcox [1983] and Rubinstein [1988]). The implication for comparing "normal" states of human security threat responses systems with these same systems in "hyper-alarmed" states is obvious.

D. Heterogeneity Across Domains for the Same Individuals

Finally, one can imagine that preferences might vary for the same individual across decision domains. Someone might be slightly risk averse over financial decisions, but exhibit risk loving behavior with respect to health decisions involving risk. In some economic studies, when analysts assume the existence of "perfect markets" between commodities defined across these domains, such cross-domain heterogeneity is ruled out *a priori*. If the individual has a known market price at which health risks can be traded off against financial risks, then one can formally model this as if there is one composite wealth, effectively assuming that there is just one risk attitude to all components of that aggregate, composite wealth. In fact, this is generally an implausible assumption. When relaxed it leads to a venerable but neglected literature on multi-attribute and multivariate risk attitudes, reviewed by Andersen, Harrison, Lau and Rutström [2011]. Consequently, social and behavioral scientists investigating risk responses who are unaware of this literature may rashly dismiss the relevance of economists' models and deprive themselves of the tools with offer the best inferential power where precise estimation is concerned.

One can model interactions between these varieties of heterogeneity. Hence a DM might behave consistently with an EUT model over financial decisions and an RDU model over health decisions, and have state-dependent preferences in each domain (e.g., after marrying or having children). In investigating the human risk response system under the stress of politically motivated hyper-alarm, we should want to know whether people are being rendered more risk averse, or more pessimistic about probabilities for more favorable outcomes, or are being triggered to manifest loss aversion. Let us now turn to directly consider these theoretical elements and their measurement.

3. Heterogeneity in Risk Preference Structures

Define the risk premium as the difference between the actuarial expected value (EV) of a risky prospect and the certain amount of money an individual would accept in exchange for giving it up (the "certainty equivalent," or CE). Assume there is no bargaining process causing the individual to strategically mis-state this CE if asked for it directly or indirectly. Focus initially on risks that have objective probabilities attached to outcomes. The EV does not then depend on the DM's attitudes to risk or subjective beliefs about these probabilities: it is just a matter of arithmetic.

We consider three core models of decision-making under objective risk. All three agree on the risk premium. What they disagree about is how to explain it in terms of three (virtually or mechanistically implemented) data-generating processes.

The core model is Expected Utility Theory (EUT), and posits that the risk premium is explained solely by an aversion to variability of earnings from a prospect. We say "variability" rather than just variance, because the DM can be averse to skewness or kurtosis.

The second model is Rank-Dependent Utility (RDU), and further posits that DMs may be pessimistic or optimistic with respect to the probabilities of outcomes. RDU does not rule out aversion to variability of earnings, but augments it with an additional psychological process. Both EUT and RDU assume that individuals asset integrate, in the sense that they net out framed losses from some endowment. This means that if the subject is told that there is a $100 house endowment, and that one lottery then involves a $15 gain and a $10 loss relative to that endowment with equal probability, the DM evaluates outcomes of $115 and $90.[13]

The third model is Cumulative Prospect Theory (CPT), which adds an aversion to losses as a possible pathway to the risk premium, and also adds the assumption that gross gains and losses matter because individuals do not locally asset integrate and evaluate net gains or losses. Using the above example, the CPT DM would evaluate outcomes of +$15 and -$10 if the $100 was the reference point for comparative valuation. The reference point is a "free parameter" in CPT, and must be identified on the basis of empirical investigation in any application or test of CPT.

A. Expected Utility Theory

Assume that utility of income is defined by a utility function U(x), where x is the lottery prize. Under EUT the probabilities for each outcome x_j, $p(x_j)$, are those that are induced by the experimenter, so expected utility is simply the probability weighted utility of each outcome in each lottery. Once the utility function is estimated, it is a simple matter to evaluate the implications for risk aversion. The concept of risk aversion traditionally refers to "diminishing marginal utility," which is driven by the curvature of the utility function, which is in turn given by the second derivative of the utility function. Although somewhat loose,

[13] We call this "local asset integration" because it refers to endowments in the laboratory session. If the subject also integrates wealth that they have "outside" the lab than we refer to "global asset integration."

this can be viewed as characterizing individuals that are averse to mean-preserving increases in the variance of returns.[14]

B. Rank-Dependent Utility

The RDU model of Quiggin [1982] extends the EUT model by allowing for decision weights on risky choice outcomes. These decision weights reflect probability weights on objective probabilities. The decision weights are defined after ranking the prizes that occur within prospects, from largest to smallest. The largest prize receives a decision weight equal to the weighted probability for that prize: the decision weight reflects the probability weight of getting at least that prize. The decision weight on the second largest prize is the probability weight of getting at least that second largest prize, minus the decision weight of getting the highest prize, and so on similarly for other prizes.

The Dual Theory (DT) specification of Yaari [1987] is the special case of the RDU model in which the utility function is assumed to be linear. Hence diminishing marginal utility can have no influence on the risk premium, and the only thing that can explain the risk premium is "probability pessimism," that is, a representation of the decision weights of expected gains that decline with their relative magnitudes.

---

14 The third and fourth derivatives can similarly, and loosely, be viewed as characterizing attitudes under EUT towards skewness and kurtosis. The literature on "higher-order risk preferences" (Eeckhoudt and Schlesinger [2006]; Ebert and Wlesen [2014]) develops these ideas, and the sense in which these statements can be made more precise.

C. Cumulative Prospect Theory

The key innovation of the CPT model of Tversky and Kahneman [1992], in comparison to EUT and RDU, is to allow sign-dependent preferences, where risk attitudes depend on whether the agent is evaluating a gain or a loss relative to some reference point. The concept of loss aversion, based on sign-dependent preferences, has been formalized in different ways in the literature. It captures the general notion that "losses loom larger than gains" when evaluating risky prospects with gains and losses.

Loss aversion is conventionally defined by what we call "utility loss aversion." This arises when the absolute value of the utility decrement of a unit loss is bigger than the utility increment of a unit gain.[15]

What if the decision weights for the gain domain differ from the probability weighting functions for the loss domain? There is nothing *a priori* in CPT to rule this out. Even if the basic utility functions for gains and losses are linear, and conventional utility loss aversion is absent, this could induce the same behavior as if there were utility loss aversion. This is called "probabilistic loss aversion" by Schmidt and Zank [2008, p.213]. Imagine that there is no probability weighting on the gain domain, so the decision weights are the objective probabilities, but that there is some probability weighting on the loss domain. Then one could easily have losses weighted more than gains, from the implied decision weights.

15 To make a statement like this, we need to be able to compare utility changes in the gain domain and the loss domain for the same DM. This means that we cannot just have a utility scale that allows any order-preserving transformation: otherwise one could choose utility numbers that violated the hypothesis. In turn, this means that we have to be more restrictive than allowing positive affine transformations, and limit ourselves to defining utility on a ratio scale rather than an interval scale.

D. Estimating Risk Preference Structures in Empirical Choice Data

We describe the analysis by which we determine whether an observed DM is better characterized as an EUT, an RDU, or a CPT agent. We then show, in tabular form, estimates we and co-authors have obtained of the highest likelihood mixtures of these modes of agency in various studied populations. These results constitute a current summary description of the emerging target for baseline specification of the structure of human risk preferences, as produced by processes that are presumed to include biological and cultural evolution, and social and individual learning.

Assume that utility of income reflects constant relative risk aversion (CRRA), defined by

$$U(x) = x^{(1-r)}/(1-r) \tag{1}$$

where x is a lottery prize and $r \neq 1$ is a parameter to be estimated. Then r is the coefficient of CRRA for an EUT individual: $r=0$ corresponds to risk neutrality, $r<0$ to a risk loving attitude, and $r>0$ to risk aversion, as defined earlier.

Let there be J possible outcomes in a lottery defined over objective probabilities. Under EUT the probabilities for each outcome x_j, $p(x_j)$, are those induced by the experimenter, so expected utility (EU) is simply the probability weighted utility of each outcome in each lottery i:

$$EU_i = \sum_{j=1,J} [\, p(x_j) \times U(x_j) \,]. \tag{2}$$

The original RDU model of Quiggin [1982] extends the EUT model by allowing for decision weights on lottery outcomes. The specification of the utility function is the same parametric specification (1) considered for EUT.[16] To calculate decision weights under RDU one replaces expected utility defined by (2) with RDU:

$$RDU_i = \sum_{j=1,J} [\ \omega(p(x_j)) \times U(x_j)\ ] = \sum_{j=1,J} [\ \omega_j \times U(x_j)\ ] \tag{3}$$

where

$$\omega_j = \omega(p_j + \ldots + p_J) - \omega(p_{j+1} + \ldots + p_J) \tag{4a}$$

for j=1,... , J-1, and

$$\omega_j = \omega(p_j) \tag{4b}$$

for j=J, with the subscript j ranking outcomes from worst to best, and ω(p) is some probability weighting function.

We consider three popular probability weighting functions. The first is the "power" probability weighting function considered by Quiggin [1982], with curvature parameter γ:

$$\gamma(p) = p^{\gamma} \tag{5}$$

So $\gamma \neq 1$ is consistent with a deviation from the conventional EUT representation. Convexity of the probability weighting function, when $\gamma>1$, is said to reflect "pessimism" and generates, if one assumes, for simplicity, a "linear" utility function, a risk premium since $\omega(p) < p$ for all p and hence the RDU expected value weighted by ω(p) instead of p has to be less than the EV weighted by p.

[16] To ease complexity of notation we use the same parameter r because the context always make it clear if this refers to an EUT or RDU model.

The second probability weighting function is the "inverse-S" function popularized by Tversky and Kahneman [1992]:

$$\omega(p) = p^{\Upsilon} / (p^{g} + (1-p)^{\gamma})^{1/Y} \qquad (6)$$

This function exhibits inverse-S probability weighting (optimism for small p, and pessimism for large p) for $\gamma<1$, and S-shaped probability weighting (pessimism for small p, and optimism for large p) for $\gamma>1$.

The third probability weighting function is a general functional form proposed by Prelec [1998] that exhibits considerable flexibility. This function is

$$\omega(p) = \exp\{-\eta(-\ln \varphi)^{j}\}, \qquad (7)$$

and is defined for $0<p\leq1$, $\eta>0$ and $\varphi>0$.

[Table 1 about here]

In the experiments compared in Table 1 evaluating EUT and RDU, subjects made K binary choices between lotteries defined in the gain frame, where K typically ranged between 30 and 100. After all decisions were made one of the K choices was chosen at random to be played out in accordance with the choices of the subject. Under EUT this experimental payment protocol provides incentives for truthful binary choices.[17] The batteries of lottery pairs used were carefully selected for the purpose of identifying whether any given subject behaves more consistently under EUT or under RDU.

[17] Harrison and Swarthout [2014] discuss the evidence for this experimental payment protocol, particularly when drawing inferences about RDU models. Their findings just make our classifications of subjects as EUT or RDU more conservative with respect to EUT (i.e., we are more likely with this payment protocol to classify subjects as RDU than if the protocol had no effect).

To evaluate RDU preferences we estimate an RDU model for each individual. We consider the CRRA utility function (1) and one of three possible probability weighting functions defined earlier by (5), (6) and (7). For our purposes of classifying subjects as EUT or RDU it does not matter which of these probability weighting functions characterize behavior: the only issue here is at what statistical confidence level we can reject the EUT hypothesis that $\omega(p) = p$.

Perusal of Table 1 shows that both EUT and RDU preference structures are very common across a range of populations. These include subjects in the USA and in developing and poor countries. They include substantial portions of men and women, undergraduate students and adult subjects drawn from broad community samples, and subjects of varying levels of income, wealth, and literacy. Not shown in the table is the fact that the overwhelming majority of subjects in all of the samples were observed to be moderately risk averse.

Table 1 shows observations of estimated proportions of CPT agents for a minority of our studies. One can only include CPT in a mixture model for data that included choices in which prospect contexts include outcomes framed by subjects as net losses, and where some prospects used a "mixed frame" in which some outcomes were gains and some outcomes were losses. This involves design and logistical complications in experiments that are not always worth paying, because in our experience the increment to mixture model estimation precision added by including CPT has typically been small as long as RDU is included in the model. This methodological judgment derived from experience should perhaps be put on hold, and re-evaluated, when researchers turn to a fundamentally new phenomenal target, such as the response of security-threat systems under hyper-alert stress.

Tversky and Kahneman [1992; p. 309] popularized the functional forms we often see for loss aversion, using a CRRA specification of utility:

$$U(m) = m^{1-\alpha} / (1-\alpha) \text{ when } m \geq 0 \tag{8a}$$

$$U(m) = -\lambda[(-m)^{1-\beta} / (1-\beta)] \text{ when } m < 0, \tag{8b}$$

and where λ is the loss aversion parameter. Here we have the assumption that the degree of loss aversion for small unit changes is the same as the degree of loss aversion for large unit changes: the same λ applies locally to gains and losses of the same monetary magnitude around 0 as it does globally to any size gain or loss of the same magnitude.

We allow flexibility in the probability weighting for losses and gains with the power probability weighting function by using

$$\omega(p) = p^{\omega+} \text{ for } m \geq 0 \tag{9a}$$

$$\omega(p) = p^{\omega-} \text{ for } m < 0 \tag{9b}$$

and where the p in question is the objective probability associated with that specific m. For the inverse-S function we use

$$\omega(p) = p^{\gamma+} / (p^{\gamma+} + (1-p)^{\gamma+})^{1/\gamma+} \text{ for } m \geq 0 \tag{10a}$$

$$\omega(p) = p^{\gamma-} / (p^{\gamma-} + (1-p)^{\gamma-})^{1/\gamma-} \text{ for } m < 0. \tag{10b}$$

For the Prelec function we use

$$\omega(p) = \exp\{ -\eta^{+} (-\ln p)^{\varphi+} \} \tag{11a}$$

$$\omega(p) = \exp\{ -\eta^{-} (-\ln p)^{\varphi-} \}. \tag{11b}$$

It is important to be explicit about these specifications of the probability weighting functions across gain and loss frames, since it is a key, but neglected, component of CPT.

We do not define the three models in terms of specific parameters. There is a tradition that defines "the CPT model" by patterns of risk attitudes, such as the "fourfold pattern of risk aversion." We do not want to constrain the model *a priori* in this manner, and in fact many of the alleged empirical regularities are simply false (e.g., Wilcox [2015]). We see no reason to reject a model because the empirical folklore claimed specific parameters and those parameter values ended up being false. Evolutionary psychologists should be alerted to this, since the majority of psychologists who have incorporated formal versions of prospect theory into their models have assumed a specifically parameterized form of CPT. In the research program to model the effects of hyper-activating threat signals, we recommend that this be avoided.

The final column of Table 1, applying to two of the samples, is of interest. It reflects use of a mixture model that combines the "aspiration" (A) model of a psychologist (Lopes [1995]) with a generalized RDU model. Lopes suggested a model focusing on a specific member of the RDU family of models, but we generalize. A is a dual-criteria decision model, like many psychological ones, including the original, pre-cumulative, version of prospect theory due to Kahneman and Tversky [1979], which also features in Table 1 as "OPT." In Lopes's model, the decision weights in SP/A theory derive from a process by which the DM attends to both a prospect's achievement of a minimum security level *and* lower-probability but especially attractive upside potential. The modeled inverted- S shape function reflects higher attention weighting of particularly bad outcomes, so as to take special care to avoid them, while simultaneously over-weighting the very best outcomes. In effect, a DM with this weighting function is biased in favor of security but attends to especially attractive prospects as long as the security threat from doing so is below a threshold value. It will be noted that this model performs extremely powerfully in the mixture models in which it was included. In fact the best performing model of all reported by Andersen, Harrison, Lau, and Rutström [2014] mixes EUT and (RDU + A) as the "top" criterion, and then nests the RDU / A mixture at a second level of choice selection.

The column of Table 1 headed "DA" shows the unimpressive performance of a "disappointment aversion" model due to Gul [1979], according to which DMs evaluate prospects according to an augmented version of EUT, in which they also take into account the extent to which an outcome differs from the CE of the prospect.

4. Heterogeneity in Subjective Beliefs About Probabilities

The notion that subjective probabilities can be usefully viewed as prices at which one might trade has been a common one in statistics, and is associated with de Finetti [1937][1970] and Savage [1971]. It is also standardly assumed in the vast literature on gambling, particularly on the setting of odds by bookies and parimutual markets (Epstein [1977, p. 298ff.]). The central insight is that subjective probabilities of events are marginal rates of substitution between contingent claims, where the contingencies are events to which the probabilities refer. There are then myriad ways in which one can operationalize this notion of a marginal rate of substitution.

Scoring rules are procedures that convert a "report" by an individual into a lottery defined over the outcome of some event. The formal link between scoring rules and optimizing decisions by agents is also familiar, particularly in Savage [1971]. A scoring rule is simply a way of translating a reported probability or belief into some earnings based on the actual outcome. It is convenient to think about scoring rules for binary events, and then scoring rules for continuous events. An example of a binary event might be "A typical American male will live 70 or more years. True or false?" The corresponding continuous event would be "How long will the typical American male live?"

A. Scoring Rules for Binary Events

A scoring rule for binary events asks the subject to make some report θ, and then defines how an elicitor pays a subject depending on their report and the outcome of the event. This framework for eliciting subjective probabilities can be formally viewed from the perspective of a trading game between two agents: you give me a report, and I agree to pay you $X if one outcome occurs and $Y if the other outcome occurs. The scoring rule defines the terms of the exchange quantitatively, explaining how the elicitor converts the report from the subject into a lottery. We use the terminology "report" because we want to view this formally as a mechanism, and do not want to presume that the report is in fact the subjective probability π of the subject. In general, it is not.

The Quadratic Scoring Rule (QSR) for binary events was apparently first used in incentivized experiments by McKelvey and Page [1990], and is by far the most popular. In most early applications the subject was implicitly or explicitly assumed to be risk-neutral, and we return to this issue below.[18] There are many other "proper" scoring rules in use.

[18] McKelvey and Page [1990] augmented the scoring rule procedure with a "binary lottery" payment procedure to induce risk-neutrality. In theory the subject earns "points" in the scoring rule, which convert in a linear manner into an increased probability of winning some lottery defined over a high prize and a low prize. There is considerable controversy over the behavioral validity of this procedure, reviewed in Harrison, Martínez-Correa and Swarthout [2013].

The QSR can be defined in terms of two positive parameters, α and β, that determine a fixed reward the subject gets and a penalty for error, respectively. Assume that the possible outcomes are A or B, where B is the complement of A, that θ is the reported probability for A, and that Θ is the true binary-valued outcome for A. Hence $\Theta=1$ if A occurs, and $\Theta=0$ if it does not occur (and thus B occurs instead). The subject is paid $S(\theta|A$ occurs$) = \alpha - \beta(\Theta-\theta)^2 = \alpha - \beta(1-\theta)^2$ if event A occurs and $S(\theta|B$ occurs$) = \alpha - \beta(\Theta-\theta)^2 = \alpha - \beta(0-\theta)^2$ if B occurs. In effect, the score or payment penalizes the subject by the squared deviation of the report from the true binary-valued outcome, Θ, which is 1 and 0 respectively for A and B occurring. An omniscient seer would obviously set $\theta = \Theta$. The fixed reward is a convenience to ensure that subjects are willing to play this trading game, and the penalty function defines the penalty from not being an omniscient seer. In the experiments reported in Andersen, Fountain, Harrison and Rutström [2014] $\alpha = \beta = \$100$, so subjects could earn up to \$100 or as little as \$0. If they reported 1 they earn \$100 if event A occurs or \$0 if event B occurs; if they reported ¾ they earn \$93.75 or \$43.75; and if they reported ½ they earn \$75 no matter what event occurs.

It is intuitively obvious, and also well known in the literature (e.g., Winkler and Murphy [1970] and Kadane and Winkler [1988]), that risk attitudes will affect the incentive to report one's subjective probability "truthfully" in the QSR.[19] A sufficiently risk averse agent is drawn to a report of ½, and varying degrees of risk aversion will cause varying

[19] There exist mechanisms that will elicit subjective probabilities without requiring that one correct for risk attitudes, such as the procedures proposed by Köszegi and Rabin [2008; p.199], Karni [2009], Grether [1992], Holt and Smith [2009], Offerman, Sonnemans, van de Kuilen and Wakker [2009] and Hao and Hauser [2012], discussed further below. The last four employ these mechanisms in an experimental evaluation. All of these elicitation procedures have some potential problems. The first is the poor incentive properties around the true subjective belief. This is particularly the case when using the Becker, DeGroot and Marschak [1964] procedures, where mis-reporting leads to very small expected lost earnings. The second is that explanations of these procedures are not easy for subjects to understand.

distortions in reports from subjective probabilities. If we knew the form of the (well-behaved) utility function of the subjects, and their degree of risk aversion, we could infer back from any report what subjective probability they must have had. This is exactly what Andersen, Fountain, Harrison and Rutström [2014] illustrate, recognizing that one only ever has *estimates* of the true degree of risk aversion.

B. Scoring Rules for Continuous Events

The DM in this setting reports her subjective beliefs in a discrete version of a QSR for continuous distributions, developed by Matheson and Winkler [1976]. Partition the domain into K intervals, and denote as r_k the report of the density in interval $k = 1,\ldots, K$. Assume for the moment that the DM is risk neutral, makes decisions consistently with SEU, and that the full report consists of a series of reports for each interval, $\{ r_1, r_2,\ldots, r_k ,\ldots, r_K \}$ such that $r_k \geq 0 \ \forall k$ and $\sum_{i = 1,\ldots, K} (r_i) = 1$. Figure 1 illustrates the case in which $K = 10$.

If k is the interval in which the true value lies, then the payoff score is from Matheson and Winkler [1976; p.1088, equation (6)]:

$$S = (2 \times r_k) - \sum_{i = 1, \ldots,K} (r_i)^2 \tag{12}$$

The reward in the score is a doubling of the report allocated to the true interval, and a penalty that depends on how these reports are distributed across the K intervals. The subject is rewarded for accuracy, but if that accuracy misses the true interval the punishment is severe. The punishment includes all possible reports, including the correct one. Let reports consist of 100 tokens to be allocated across the intervals.

To ensure complete generality, and avoid any DM facing losses, again allow some endowment, α, and scaling of the score, β. We then get the generalized scoring rule

$$\alpha + \beta \left[(2 \times r_k) - \sum_{i=1,\ldots,K} (r_i)^2 \right] \tag{13}$$

where we initially assumed α=0 and β=1 in (12). We can assume α>0 and $\beta\neq 0$ to get the payoffs to any positive level and units we want. In our elicitation procedures $K = 10$, as shown in Figures 1 and 2. The QSR we use in the behavior reported later, and underlying the displays in Figures 1 and 2, uses $\alpha = \beta = 25$. Hence the maximum payoff possible, if all tokens are allocated to one interval, is $50.

[Place Figure 1 about here]

[Place Figure 2 about here]

In our application each individual selects an allocation of 100 tokens by sliding a bar for each bin, with the "histogram" representation changing in real time. Only when 100 tokens have been allocated can the allocation be submitted, and even then there is a need to actively confirm the choice. This design extends the binary QSR interface single-slider developed by Andersen, Fountain, Harrison and Rutström [2014], which allows the experimenter to use a specific QSR to generate the implied allocations without burdening the individual with messy formulae. The allocation is always initialized at 0 tokens for every interval.

C. Controlling for Risk Attitudes

We do not know whether the subject whose subjective beliefs about probabilities we are estimating is risk neutral. Indeed, the weight of evidence from past laboratory and field experiments clearly suggests that subjects will be modestly risk averse over the prizes they face. One approach is to try to induce risk neutrality using binary lottery procedures, as reviewed in Harrison, Martínez-Correa, Swarthout and Ulm [2014][2015] and applied in Harrison and Phillips [2014]. Another approach is to estimate risk attitudes and consider the extent to which they distort reported beliefs. We favor the latter approach in general, since it puts less strain on the subject understanding the logic of the binary lottery procedure.

Subjective Probabilities over Binary Events

It is well-known that risk aversion can significantly affect inferences from applications of the QSR to eliciting subjective *probabilities* over *binary* events (Winkler and Murphy [1970], Kadane and Winkler [1988]). Andersen, Fountain, Harrison and Rutström [2014] demonstrate how one can estimate utility functions and/or probability weighting functions using choices over lotteries with objective probabilities, assume that the same functions apply to scoring rule responses, and estimate latent subjective probabilities.

Subjective Beliefs over Continuous Events

Harrison, Martínez-Correa, Swarthout and Ulm [2012] characterize the implications of the general case of a risk averse agent when facing the QSR and reporting subjective *distributions* over *continuous* events, and find, remarkably, that these concerns do not apply with anything like the same force under Subjective Expected Utility (SEU). For empirically plausible levels of risk aversion, one can reliably elicit the most important features of the latent subjective belief distribution for an SEU agent without undertaking calibration for risk attitudes. Providing that our subjects exhibit the modest levels of risk aversion found universally in the lab and field settings for stakes of the levels typically used in economic experiments (e.g., Harrison and Rutström [2008]), and that we are willing to assume SEU for the individual, these results provide the basis for using the reported distributions as if they are the true, subjective belief distributions.

Recovering Subjective Belief Distributions

A maintained assumption in these comforting results is that the DM behaves consistently with SEU. Harrison and Ulm [2015] generalize these findings further, showing how one can exactly calibrate latent subjective belief distributions for individuals that behave consistently with EUT or RDU when faced with choices over objective lotteries.

Recovered Beliefs Are Conditional on the Assumed Model of Risk Preferences

Figure 3 illustrates the importance of determining the correct model of risk preferences, a theme that goes back to Savage [1971][1972]: one cannot recover subjective beliefs without making some assumptions about the underlying model of risk preferences. These are the elicited and recovered beliefs of one person about the fraction of deaths due to cancer that the Center for Disease Control attributes to smoking. This subject is a young, female smoker, so these are important beliefs for her health decisions. If we use a 1% significance level in testing for violation of EUT then we characterize her as an EUT DM, and her recovered beliefs closely track the reported beliefs. However, if we use a 5% or higher significance level, then we characterize her as an RDU DM, and recover latent subjective beliefs that are much closer to the true facts. She underestimates the risk of smoking no matter how we characterize her risk preferences, but the misperception is clearly greater if we model her as an EUT DM. So from a *qualitative* perspective we do not need to know whether her risk preferences are EUT or RDU, but to ascertain the *quantitative magnitude* of her misperception we do need to correctly know those preferences.

[Place Figure 3 about here]

5. Conclusion

Human dispositions to choose under conditions of risk and uncertainty involve a number of structural elements that can be independently elicited and jointly estimated. We reviewed methods for estimating first-order risk preferences, risk preference structure (i.e., EUT or alternatives), and subjective beliefs about probabilities. Other elements, second-order risk preferences, were mentioned but not reviewed. Modeling choice under uncertainty (or ambiguity) involves further methodological complications. We focused on observed heterogeneity, along multiple dimensions, in these aspects of risk response structure for two reasons. First, economists often model risk in ways that foreclose observation of heterogeneity. This can lead psychologists and other behavioral scientists to doubt that economists' theoretical approach is sufficient empirically sensitive at a grain of analysis that engages with real behavior at the relevant level of abstraction. Second, in the specific context of studying evolved biases in human risk response, when psychologists *have* borrowed economists' models for the sake of specification rigor, they have sometimes been led by the imposition of horse race methods (which are one strategy for imposing homogeneity) into thinking that evolution favored a uniform bias, for example loss aversion. We presented evidence against this hypothesis, at least insofar as currently operative risky choice patterns are taken to be evidence for biases inherited from human ancestors. The target model for evolutionary modeling of the baseline risky choice structure, according to the evidence presented, is one in which optimization features prominently in data, but not more prominently than models according to which tail probabilities are weighted pessimistically in subjective expectations. Another explanatory target for evolutionary psychologists is the recurrent finding that most people are moderately risk averse.

Concerning the vulnerability of human risk response system to exploitation of dispositions to hyper-alarm suggested by Woody and Szechtman [2016], our findings enrich the range of available structural specifications of hypotheses. Exaggeration of loss aversion does not appear to be recommended as the first place to look, contrary to what might be gathered from popular accounts such as Ariely [2008]. One possibility is that risk aversion increases in domains of preference that scare-mongers emphasize. A more likely first suspect, we suggest, is distortions of probability weighting functions, again in specific domains. This suggestion raises intriguing further questions in advance of the research. Are people who are disposed to EUT-consistent choice in pre-alarmed conditions less prone to exploitation by scare-mongers than RDU choosers? Or, given the strong performance in two mixture models of Lopes's dual criteria model, are some people vulnerable to having their reflective processes suppressed, leading to runaway dominance of processes that focus exclusively on security levels? As we have documented, experimental economists are equipped with a set of sharp investigative and analytical tools with which to help psychologists test mathematically well specified formulations of these and related hypotheses.

Table 1: Observed proportions of EUT, RDU, and CPT DMs based on lottery-choice experiments

					Number of…		Percent Behaving Consistently with…						
Study	**Country**	**Sample**	**Treatment**	**Econometrics**	**Subjects**	**Choices**	**EUT**	**DT**	**RDU**	**DA**	**OPT**	**CPT**	**A**
Harrison and Rutström [2009]	USA	Undergraduates		Pooled mixture model	158	9311	55%				45%		
Harrison, Humphrey and Verschoor [2010]	Ethiopia, India and Uganda	Peasant farmers		Pooled mixture model	531	4248	46%		54%				
Harrison and Ng [2016]	USA	Undergraduates		Individual estimation	102	8160	49%		51%				
Harrison and Swarthout [2016]	USA	Undergraduates	House money	Individual estimation	177	17700	9%	0%	76%	0%		15%	
		Undergraduates	Earned endowment	Individual estimation	58	5800	7%	0%	55%	5%		33%	
		MBA students	House money	Individual estimation	94	9400	26%	0%	61%	0%		13%	
Harrison and Ulm [2016]	USA	Undergraduates		Individual estimation	65	3250	33%		67%				
Harrison, Ng, Swarthout and Ulm [2016]	USA	Undergraduates		Individual estimation	76	11020	60%		40%				
Harrison and Ross [2014]	South Africa	University faculty		Individual estimation	193	9650	57%		43%				
Andersen, Harrison, Lau and Rutström [2010]	USA	Undergraduates	Deal Or No Deal	Pooled mixture model	125	870			7%				93%
Andersen, Harrison, Lau and Rutström [2014]	U.K.	Game show contestants	Deal Or No Deal	Pooled mixture model	461	2317			35%				65%

Figure 1: Belief Elicitation Interface

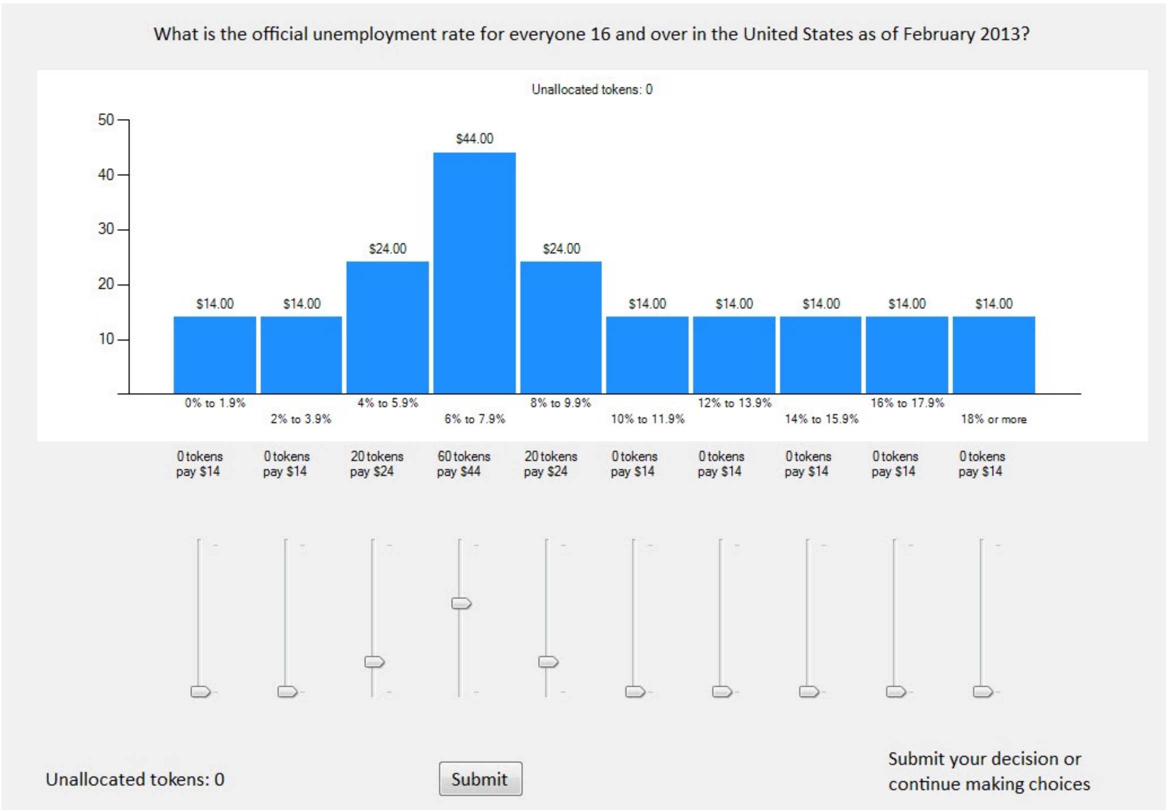

Figure 2: Possible Belief Elicitation Response

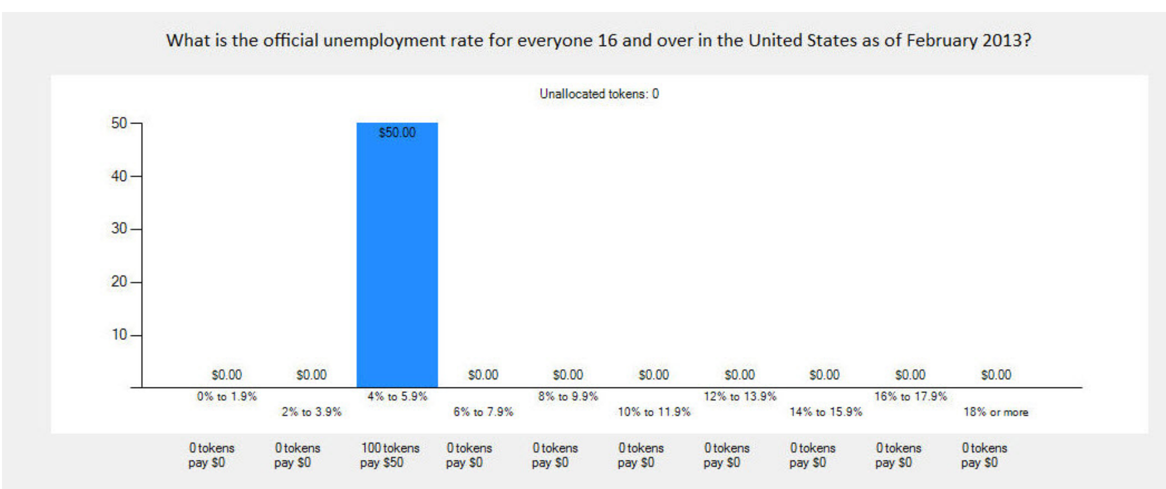

Figure 3: Importance of Modeling Risk Preferences Correctly for a Young, Female Smoker

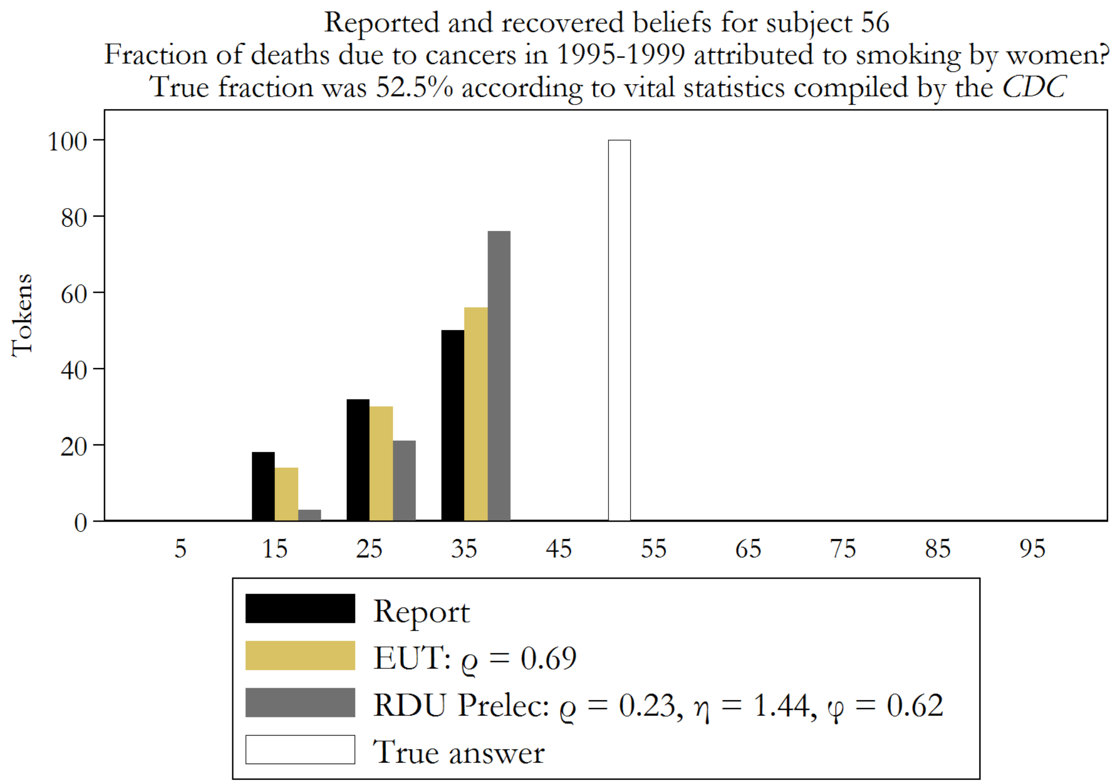

References

Ainslie, George [1992]. *Picoeconomics*. Cambridge: Cambridge University Press.

Andersen, Steffen; Fountain, John; Harrison, Glenn W., and Rutström, E. Elisabet [2014]. "Estimating Subjective Probabilities." *Journal of Risk & Uncertainty,* 48: 207-229.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I.; and Rutström, E. Elisabet [2008]. "Lost in State Space: Are Preferences Stable?" *International Economic Review*, 49[3]: 1091-1112.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten, and Rutström, E. Elisabet [2010]. "Behavioral Econometrics for Psychologists." *Journal of Economic Psychology,* 31: 553-576.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet [2011]. "Multiattribute Utility Theory, Intertemporal Utility and Correlation Aversion." *Working Paper 2011-04*. Center for Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten Igel, and Rutström, E. Elisabet [2014]. "Dual Criteria Decisions." *Journal of Economic Psychology*, 41: 101-113.

Ariely, Dan [2008]. *Predictably Irrational*. New York: Harper Collins.

Ballinger, T. Parker, and Wilcox, Nathaniel, [1997]."Decisions, Error and Heterogeneity." *Economic Journal*, 107: 1090-1105.

Becker, Gordon M.; DeGroot, Morris H., and Marschak, Jacob [1964]. "Measuring Utility By A Single-Response Sequential Method." *Behavioral Science*, 9: 226-232.

Binmore, Ken [2009]. *Rational Decisions.* Princeton: Princeton University Press.

Clark, Andy [1989]. *Microcognition.* Cambridge, MA: MIT Press.

Clark, Andy [1997]. *Being There*. Cambridge, MA: MIT Press.

de Finetti, Bruno [1937/1980]. "La Prévision: Ses Lois Logiques, Ses Sources Subjectives," *Annales de l'Institut Henri Poincairé,* 7: 1-68. English translation as "Foresight: Its Logical Laws, Its Subjective Sources." In H.E. Kyburg and H.E. Smokler (eds.) [Second Edition], *Studies in Subjective Probability.* Huntington, NY: Robert E. Krieger.

de Finetti, Bruno [1970]. "Logical Foundations and Measurement of Subjective Probability." *Acta Psychologica*, 34: 129-145.

Deaton, Angus [2010]. "Instruments, Randomization, and Learning About Development." *Journal of Economic Literature,* 48: 424-455.

Dennett, Daniel [1987]. *The Intentional Stance*. Cambridge, MA: MIT Press.

Ebert, Sebastian, and Wlesen, Daniel [2014]. "Joint Measurement of Risk Aversion, Prudence, and Temperance." *Journal of Risk and Uncertainty,* 48: 231-252.

Eeckhoudt, Louis, and Schlesinger, Harris [2006]. "Putting Risk in its Place." *American Economic Review,* 96: 280-289.

Epstein, Richard A. [1977]. *The Theory of Gambling and Statistical Logic.* San Diego: Academic Press.

Gollier, Christian [2001]. *The Economics of Risk and Time.* Cambridge, MA: MIT Press.

Grether, David M., and Plott, Charles R. [1979]. "Economic Theory of Choice and the Preference Reversal Phenomenon." *American Economic Review*, 69: 623-648.

Gul, Faruk [1991]. "A Theory of Disappointment Aversion." *Econometrica*, 59[3]: 667-686.

Harrison, Glenn W.; Humphrey, Steven J., and Verschoor, Arjan [2010]. "Choice under Uncertainty: Evidence from Ethiopia, India and Uganda." *Economic Journal,* 120: 80-104.

Harrison, Glenn W., and Lau, Morten I.[2014]. "Risk Attitudes, Sample Selection and Attrition in a Longitudinal Field Experiment." *Working Paper 2014-04*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet [2007]. "Estimating Risk Attitudes in Denmark: A Field Experiment." *Scandinavian Journal of Economics*, 109[2]: 341-368.

Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet [2009]. "Risk Attitudes, Randomization to Treatment, and Self-selection into Experiments." *Journal of Economic Behavior and Organization,* 70: 498-507.

Harrison, Glenn W.; Martínez-Correa, Jimmy; Ng, Jia Min, and Swarthout, J. Todd [2016]. "Evaluating the Welfare of Index Insurance." *Working Paper 2016-07.* Center for Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W.; Martínez-Correa, Jimmy, and Swarthout, J. Todd [2014]. "Eliciting Subjective Probabilities with Binary Lotteries." *Journal of Economic Behavior & Organization*, 101: 128-140.

Harrison, Glenn W.; Martínez-Correa, Jimmy; Swarthout, J. Todd, and Ulm, Eric [2012]. "Scoring Rules for Subjective Probability Distributions." *Working Paper 2012-10*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W., Martínez-Correa, Jimmy; Swarthout, J. Todd, and Ulm, Eric R. [2015]. "Eliciting Subjective Probability Distributions with Binary Lotteries." *Economics Letters*, 127: 68-71.

Harrison, Glenn W., and Ng, Jia Min [2016]. "Evaluating the Expected Welfare Gain From Insurance." *Journal of Risk and Insurance,* 83: 91-120.

Harrison, Glenn W., and Phillips, Richard D. [2014]. "Subjective Beliefs and the Statistical Forecasts of Financial Risks: the Chief Risk Officer Project." In T.J. Andersen (ed.) *Contemporary Challenges in Risk Management.* New York, Palgrave Macmillan.

Harrison, Glenn W., and Ross, Don [2016]. "Varieties of Paternalism and the Heterogeneity of Utility Structures." *Working Paper 2016-06.* Center for Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W., and Rutström, E. Elisabet, [2008]. "Risk Aversion in the Laboratory." In J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments,* Volume 12. Bingley, UK: Emerald, Research in Experimental Economics.

Harrison, Glenn W., and Rutström, E. Elisabet [2009]. "Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral." *Experimental Economics,* 12: 133-158.

Harrison, Glenn W., and Swarthout, J. Todd [2014]. "Experimental Payment Protocols and the Bipolar Behaviorist." *Theory and Decision,* 77: 423-438.

Harrison, Glenn W., and Swarthout, J. Todd [2016]. "Cumulative Prospect Theory in the Laboratory: A Reconsideration." *Working Paper 2016-05.* Center for Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W., and Ulm, Eric R. [2015]. "Recovering Subjective Probability Distributions." *Working Paper 2015-01*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Harrison, Glenn W., and Ulm, Eric [2016]. "Recovering Subjective Probability Distributions." *Working Paper 2016-01.* Center for Economic Analysis of Risk, Robinson College of

Business, Georgia State University.

Heckman, James J. [1992]. "Randomization and Social Program Evaluation." In Charles Manski and Irwin Garfinkel (eds.), *Evaluating Welfare and Training Programs.* Cambridge, MA: Harvard University Press.

Hey, John D., and Orme, Chris [1994]. "Investigating Generalizations of Expected Utility Theory Using Experimental Data." *Econometrica*, 62[6]: 1291-1326.

Holt, Charles A., and Smith, Angela M. [2009]. "An Update on Bayesian Updating." *Journal of Economic Behavior & Organization*, 69: 125-134.

Hutchins, Edwin [1995]. *Cognition in the Wild.* Cambridge, MA: MIT Press.

Kadane, J. B. and Winkler, Robert L. [1988]. "Separating Probability Elicitation from Utilities." *Journal of the American Statistical Association*, 83[402]: 357-363.

Kagel, John; Battalio, Raymond, and Green, Leonard, eds. [2007]. *Economic Choice Theory: An Experimental Analysis of Animal Behavior.* Cambridge: Cambridge University Press.

Kahneman, Daniel, and Tversky, Amos [1979]. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica,* 47: 263-292.

Karni, Edi [2009]. "A Mechanism for Eliciting Probabilities." *Econometrica*, 77[2]: 603-606.

Keane, Michael P. [2010a]. "Structural vs Atheoretic Approaches to Econometrics." *Journal of Econometrics,* 156: 3-20.

Keane, Michael P. [2010b]. "A Structuralist Perspective on the Experimentalist School." *Journal of Economic Perspectives,* 24: 47-58.

Köszegi, Botond, and Rabin, Matthew [2008]. "Revealed Mistakes and Revealed Preferences." In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics: A Handbook,* New York: Oxford University Press.

Leamer, Edward E. [2010]. "Tantalus on the Road to Asymptopia." *Journal of Economic Perspectives,* 24: 31-46.

Lopes, Lola L. [1984]. "Risk and Distributional Inequality." *Journal of Experimental Psychology: Human Perception and Performance*, 10[4]: 465-484.

Matheson, James E., and Winkler, Robert L. [1976]. "Scoring Rules for Continuous Probability Distributions." *Management Science*, 22[10]: 1087-1096.

McKelvey, Richard D., and Page, Talbot [1990]. "Public and Private Information: An Experimental Study of Information Pooling." *Econometrica*, 58[6]: 1321-1339.

Offerman, Theo; Sonnemans, Joep; van de Kuilen, Gijs, and Wakker, Peter P. [2009]. "A Truth-Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes." *Review of Economic Studies*, 76[4]: 1461-1489.

Prelec, Drazen [1998]. "The Probability Weighting Function." *Econometrica,* 66: 95-113.

Quiggin, John [1982]. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization,* 3: 323-343.

Rubinstein, Ariel [1988]. "Similarity and Decision Making under Risk [Is There a Utility Theory Resolution of the Allais Pradox?]." *Journal of Economic Theory*, 46: 145–153.

Savage, Leonard J. [1954]. *The Foundations of Statistics.* New York: Wiley.

Savage, Leonard J. [1971]. "Elicitation of Personal Probabilities and Expectations." *Journal of American Statistical Association*, 66: 783-801.

Savage, Leonard J. [1972]. *The Foundations of Statistics* [Second Edition]. New York: Dover Publications.

Schmidt, Uri, and Zank, Horst [2008]. "Risk Aversion in Cumulative Prospect Theory." *Management Science,* 54: 208–216.

Stigler, George J., and Becker, Gary S. [1977]. "De Gustibus Non Est Disputandum." *American Economic Review*, 67[2]: 76-90.

Tversky, Amos, and Kahneman, Daniel [1992]. "Advances in Prospect Theory: Cumulative Representations of Uncertainty." *Journal of Risk and Uncertainty,* 5: 297-323.

U.S. Department of Health and Human Services [2004]. *The Health Consequences of Smoking: A Report of the Surgeon General.* Washington, D.C.: Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.

von Neumann, John, and Mogenstern, Oscar [1944]. *Theory of Games and Economic Behavior.* Princeton: Princeton University Press.

Wilcox, Nathaniel [1983]. "Lottery Choice: Incentives, Complexity, and Decision Time." *Economic Journal*, 103: 1397-1417.

Wilcox, Nathaniel [2015]. "Unusual Estimates of Probability Weighting Functions." *Working Paper.* Orange County, CA: Economic Science Institute, Chapman University.

Winkler, Robert L., and Murphy, Allan H. [1970]. "Nonlinear Utility and the Probability Score." *Journal of Applied Meteorology*, 9: 143-148.

Woody, Eric, and Szechtman, Henry [2016]. "Unintended Consequences of Security Motivation in the Age of the Internet: Impacts on Governance and Democracy." *Journal of Cognition and Culture*, this volume.

Yaari, Menahem [1987]. "The Dual Theory of Choice under Risk." *Econometrica,* 55: 95-115.