# Behavior Towards Endogenous Risk in the Laboratory

by

Glenn W. Harrison, E. Elisabet Rutström and Shabori Sen†

January 2010

ABSTRACT. We consider the effect of allowing for endogenous risk on behavior, in the simplest possible laboratory environment that allows us to identify effects on risk attitudes and subjective beliefs. We find that risk attitudes and subjective beliefs are indeed significantly affected by the possibility that choices might alter the probabilities in the lotteries being considered. Our design considers the evaluation of risk, and risk mitigating behavior, in a virtual environment that reflects the risk of property destruction from a forest fire. We jointly elicit risk attitudes, subjective beliefs and mitigating choices, and estimate a structural model that allows us to identify the effects of allowing for endogeneity of risk. We find evidence that endogenous risk settings do cause subjects to employ different subjective beliefs than they use in an exogenous risk setting, although risk attitudes appear stable across these settings.

† Department of Risk Management & Insurance, Robinson College of Business, Georgia State University, and Center for the Economic Analysis of Risk (Harrison); Department of Economics, College of Business Administration, University of Central Florida, USA (Rutström and Sen). E-mail: gharrison@gsu.edu, erutstrom@bus.ucf.edu and ssen@bus.ucf.edu. We thank the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746.

Most interesting choices in risky environments allow the individual to undertake mitigation or self-protection choices which alter the risk that is faced. In the business world, this is called risk management. Using the formal metaphor of economists, individuals get to choose which lottery to play out, within some bounds. Consider the mortality lottery, which has two known outcomes. By doing nothing, I face one set of risky outcomes. But if I engage in relatively healthy behavior the probabilities on those outcomes change, in ways that I may or may not discern well. And there might be some function relating mitigation efforts or expenditures and probability changes that I only know approximately. As we expand the example to include the morbidity lottery, my mitigation expenditures might also change the final outcomes themselves: as I spend more money on health care, my income is changed for every possible health status. This methodological problem is central to virtually every important social policy choice (e.g., towards global warming or debates over health insurance) as well as individual choice behavior (e.g., towards investment in education and health, or choice of job).

All of these concerns raise important questions about understanding behavior towards risk that is "endogenous" in this sense. To what extent do the traditional methods of analysis towards "exogenous risk" change when we allow for such mitigation behavior? This turns out to imply delicate issue of inference, as emphasized in the extant literature (e.g., Ehrlich and Becker (1972), Garen (1988), Shogren (1990), Shogren and Crocker (1991) (1994)). Do risk attitudes stay the same? Do the probabilities on final outcomes stay the same? Or do we find objective probabilities being replaced with subjective probabilities? Moreover, if there is some uncertainty about the functional relationship between mitigation efforts and probabilities of final outcomes, is it valid to assume that individuals behave as if they have a single subjective probability, rather than a density function of subjective beliefs?

In that case, do we need to allow for "uncertainty aversion" as well as risk aversion? Do individuals exhibit a preference for earlier resolution of the uncertainty about the effects of risk mitigation on subjective probabilities?

We begin the process of exploring these questions in an experimental setting in which we can focus on basics, and tease apart the confounds that arise because of the added endogeneity of risk management. We adapt the well-developed tools for structural econometric estimation of risk attitudes in an exogenous risk setting to evaluate the changes that are needed as one moves incrementally to an endogenous risk setting. Harrison and Rutström (2008) provide an extensive review of the procedures used to elicit and estimate risk attitudes in the exogenous risk setting, and we build on the toolkit reviewed there to consider what modeling issues have to be considered as one moves to model endogenous risk and risk management options. Several of these extensions are "familiar" to economists in the sense that they have arisen in different settings, but some are relatively novel. Our objective is to proceed slowly and deliberately to build a modeling bridge between the two settings.

The first extension of traditional modeling is to consider the possible role of endogeneity on risk attitudes, in the sense that attitudes towards risk might be "source dependant" on the type of stochastic process. To see why this is fundamental conceptually, imagine a mitigation activity that changes the *objective* probabilities of final outcomes in some known manner. Life is rarely that crisp, but assume this for now. If the agent behaves consistently with Expected Utility Theory (EUT), and we know the utility function of the agent, then we can immediately predict how choices will change.[1] We may not be able to

[1] The point being made here does not depend on whether one uses EUT or any alternative specification of choice under risk, such as rank-dependant utility or prospect theory. Those

qualitatively sign the changes, without more structure, but we can evaluate the expected utility (EU) of the revised lottery choices and say what the agent will, or should, do. But what if the agent also experiences a shift in risk attitudes? Then it is possible that there are changes in final outcomes that lead the analyst predict one thing if one assumes state (or source) independent risk attitudes but to predict another thing under alternative specifications.

The second extension is to consider the implications of risk mitigation on subjective probabilities. It is rare that the agent has a firm menu telling him how different mitigation actions change the objective probabilities of final outcomes, even if one could assume that the initial, no-mitigation probabilities were in fact objective. So we have to consider the elicitation of subjective probabilities. In principle, so far, this involves no radical extension of EUT, interpreted in the sense of Subjective EUT (SEUT).

The third extension is to allow for risk mitigation to generate compound, subjective lotteries over final outcomes. Even if the baseline probabilities of outcomes are subjective, it is arguable that the effect of mitigation is to add a layer of uncertainty. Of course, this layer might involve a reduction in uncertainty over final outcomes, but the methodological point is that the lottery becomes a compound lottery in which the agent has to estimate the effect of mitigation on probabilities. If we maintain a strict SEUT framework, then this extension is again straightforward, since we can use the Reduction of Compound Lotteries axiom and reduce this compound lottery to a simple lottery. But if that axiom is violated, then we admit the possibility of an additional layer of aversion to risk, or uncertainty aversion (referred to by some as "ambiguity aversion," but we prefer to avoid that label). This raises two

theories still allow one to talk about risk attitude: they simply model those attitudes as deriving from several possible sources, not just properties of the utility function.

inferential issues: how do the statistically degenerate subjective probabilities of SEU get characterized when we allow them to be non-degenerate subjective probability distributions, and how do we model uncertainty aversion? This is new terrain,[2] although to some extent our allowance for source-dependent risk attitudes captures some of what is hypothesized here, in the spirit of the early, and neglected, contribution of Smith (1969) to the Ellsberg Paradox.

We do not claim that this list exhausts the issues that have to be addressed, but it is already a serious and demanding list.[3] Our approach is to exploit the controls we have in a laboratory experiment to design a series of choice tasks that allow us to explore these issues. Following the general methodological approach of Andersen, Harrison, Lau and Rutström (2008), we present each subject with a series of different choice tasks to allow us to identify all of the behavioral parameters of interest. In their case they used a task to elicit utility functions and a task to elicit discount rates, and jointly estimated behavior in a structural model that connected the two theoretically. Similarly, Andersen, Fountain, Harrison and Rutström (2009a) used a task to elicit risk attitudes and a task to directly elicit subjective probabilities about some naturally-occurring event. Finally, Fiore, Harrison, Hughes and Rutström (2009) (FHHR) used a task to elicit risk attitudes and a task to elicit subjective beliefs in a setting of endogenous risk where the subject could pay to mitigate risk. We employ essentially the same methodology, extended to consider endogenous risk. To better understand the role of endogenous risk, we take the task of FHHR and "decompose" our analysis of behavior by additionally presenting subjects with comparable tasks in an exogenous risk setting.

[2] Experimental and econometric methods for doing so are proposed in Andersen, Fountain, Harrison and Rutström (2009b).
[3] For example, it ignores the role of insurance.

The first set of tasks are familiar choices over lotteries with objective, exogenous risk, following Holt and Laury (2002). In order to check for a framing effect from some of the later tasks being presented to the subjects as a loss from some endowment, we also include a task in which we vary the gain or loss frame with objective, exogenous risks. In these tasks the prizes are fixed across pairs of choices, but the probabilities of the good outcome vary exogenously across pairs. Moreover, the probability of the good outcome is the same in each choice: all that varies within each choice is whether one lottery is "safe" in the sense of offering intermediate rewards and the other lottery is "risky" in the sense of offering extreme rewards.

The second set of tasks vary this exogenous risk setting by allowing the subject to make choices over lotteries in which the prizes are the same in each lottery but the probabilities of the good outcome vary. So this allows the subject to make choices over lotteries in which the focus is on the *change* in the objective, exogenous probability. The first two sets of tasks allow us, in a recursive sense to be made clearer later, to estimate conventional risk attitudes.

The third set of tasks involves the direct elicitation of subjective beliefs about the risk of good or bad outcomes in a betting task. In this case the outcomes refer to whether a house that is owned by the agent is consumed by a forest fire in a Virtual Reality (VR) environment. The subject is given information, in the form of a series of naturalistic cues, about this risk; this aspect of the design closely follows FHHR, as explained later, and is needed in order to consider behavior in which *subjective* probabilities are employed. In this task the subject is simply given a series of bets, as if offered by an array of bookies, and asked to make a choice as to whether they are willing to bet on the house burning or not. By varying the "house odds" that the bookie implicitly offers in each bet, we can identify a switch point at which the subject is indifferent between betting on the bad outcome or not. Of course, this inference about subjective probabilities depends on the risk attitudes of the agent, since these are risky bets, but this is where our knowledge of the risk attitudes from the first two tasks allows us to make joint inferences about subjective beliefs and risk attitudes.

The final set of tasks replicates and extends the risk mitigation task of FHHR: the agent is able to spend money out of an endowment that will affect the subjective probability of the house burning down.[4] The optimal choice in this task depends on risk attitudes and subjective beliefs about the risk of the bad outcome from the forest fire. But we "know" these parameters from the first three sets of tasks, and can then determined if behavior in the endogenous risk setting of the fourth task is any different.[5] In effect, from the first three tasks we should be able to predict what the agent would do in the fourth task if endogenous risk has no effect on behavior. But if those predictions are off, then we can estimate which structural parameters explain the difference.

There are several procedural variants on each task, mainly to vary the domain of rewards to better estimate structural parameters, but the logic of the design is simple. The one item in our initial list that we do not consider here is an allowance for uncertainty aversion from endogenous risk. That extension involves a new series of modeling and inferential subtleties beyond the immediate scope of this analysis (see Andersen, Fountain, Harrison and Rutström (2009b)), but for which our design and analysis is a necessary precursor.

---

[4] Of course, insurance involves the expenditure of money out of an endowment that will affect the net value of the outcome, rather than the probability of the outcome.

[5] We put the word "know" in quotation marks as a reminder that we are estimating these parameters, and that there are then standard errors to be accounted for. Our full-information approach to joint estimation of all parameters accounts for this.

One central feature of our experiments is to elicit behavior in the context of experiments that are conducted in a laboratory environment, but that provide some of the naturalistic cues and context faced by decision makers in the field. Field experiments are limited in their ability to implement necessary controls on experimental conditions and a wide range of counterfactual scenarios; this is where laboratory experiments typically have great advantages. And yet conventional laboratory experiments are often artefactual by design (Harrison and List (2004)), raising questions about the validity of responses in that unusual environment.[6] Our solution is to use virtual experiments (VR) introduced by FHHR (p. 66): "A VR is an experiment set in a controlled lab-like environment, using typical lab or field participants, that generates synthetic field cues using Virtual Reality (VR) technology."

Our experimental design is borrowed from FHHR, with additional control tasks added. The subject is allotted a virtual house in the middle of the Ashley National Forest, Utah. The risk involved is that the house will burn down in a simulated forest fire. The subject faces two alternative fire policy regimes: no prescribed burning and the retention of the existing baseline risk of forest fires, and the use of prescribed burning and a reduction of the risk of forest fires. As in any natural scenario, the individual is not informed about the objective risk of fire with or without mitigation, and has to form his own subjective belief about the probability of the house burning down. However, the subject gets to experience computer simulations of several scenarios of forest fires before making his decision, and these simulations are dynamically rendered in a graphic-intensive, naturalistic manner.

[6] It is tempting to say that (good) lab experiments have the strength of internal validity at the cost of external validity, that (good) field experiments reverse this, and that the "virtual experiments" proposed here combine the best of both. But we find the terms internal and external validity confusing in some respects, for reasons not central here.

Since there are two alternative fire policies available to him, the subject needs to form a belief about two distinct events: the probability of the house burning down when no prescribed burn has been undertaken, and probability of the house burning down when a prescribed burn has been undertaken.

Section I provides a brief review of the literature, and theoretical results when risk is endogenous. It also gives an account of the few laboratory experiments that involve endogenous risk. Section II describes our experimental design, with special reference to the VR simulations that we use and the manner in which multiple tasks are used to "triangulate" and identify the structural parameters of interest. Results are presented in section III, and extensions proposed in section IV.

## I. Literature Review

### A. EU Model with Endogenous Risk

Suppose there are two states of nature: for example, ill health or good health, injury at work place or no injury, or natural disaster damaging your house or leaving it untouched. Let the probability of damage be $\rho$. Then the expected utility of an individual is given by:

EU= $[1-\rho]U(Y)+\rho U(Y-l)$

where $Y$ is initial wealth, $l$ is the loss experienced, and no mitigation is possible.

Ehrlich and Becker (1972) (EB) studied mitigation in the EUT framework.[7] EB distinguished between two forms of mitigation – self-insurance and self-protection. *Self-insurance* investment consists of expenditures made to reduce loss caused by the occurrence

[7] Indeed, in a review of the evolution of insurance economics since 1973, Loubergé (2000; p. 7) notes that EB "… were the first to propose a rigorous analysis of risk prevention. They coined the terms self-protection and self-insurance to to designate the two mechanisms and studied their relationship to "market insurance." For this reason, this paper may be seen as the first theoretical paper on risk management."

of the unwanted event (ill health, injury or disaster). Define the loss function as $L = L(l, c)$ where $c$ is self-insurance investment; $L'(c)$ is assumed to be negatively sloped, implying that loss decreases as self-insurance rises. The agent's problem now is to choose $c$ to maximize:

$$\text{EU} = (1-\rho)U(Y-c)+\rho U(Y-L(l,c)-c)$$

For self-insurance investment to be positive $-L'(c) > 1$ must hold. Here risk is exogenous.

*Self-protection* is investment made to reduce the probability of incurring any damages when the bad outcome occurs. Let $P = P(\rho, s)$ be the effective or endogenous risk function, where $\rho$ is the exogenous risk of damage occurring, $s$ is the expenditure on self-protection, and $P'(s) < 0$. In the absence of market or self-insurance the agent's problem is then to choose self-protection investment $s$ to maximize:

$$\text{EU} = [1-P(\rho,s)]U(Y-s)+P(\rho,s)U(Y-l-s)$$

where $l$ is the dollar damage caused. Self-protection is the only available mitigation alternative in our experiments.

In this model the probability of damage, conditional on the bad outcome occurring, is itself a function of the investment made by the individual to reduce risk. Allowing for self-protection makes risk endogenous, since here the risk $P$ is affected by the mitigation activity undertaken by the individual.

Two observations can be made about this model. First, the utility function $U(.)$ is independent of risk being endogenous or exogenous. An individual has the same utility *function*, whether or not risk can be mitigated, although the *level* of utility may of course differ in the two cases. This implies that if the utility function, exogenous risk level, and endogenous risk function are known, it is possible to predict the optimal level of mitigation.

Second, in the case of self investment, an increase in exogenous risk leads to an increase in the level of mitigation:

$$\partial c / \partial \rho = -[U'(y) - U'(x)(L'+1)] / (1-\rho)U''(y) + \rho U''(x) - \rho U'(x)(L'') \geq 0,$$

where $y = Y - c$ and $x = Y - L(l, c) - c$. In other words, self-insurance expenditure increases with an increase in risk. When risk is endogenous, however, it is not possible to unambiguously sign the comparative statics relation $\partial c / \partial \rho$.

### B. Empirical Implications of Endogeneity

One of the first studies to deal with endogeneity in empirical data is Garen (1988), who studied the effect of workplace risk on wages. To take his example, suppose $q_i$ and $r_i$ represent the probability of fatal and non-fatal injury, respectively, on a particular job. The wage $y_i$ is a function of the risks faced at work $y_i = y_i(q_i, r_i)$. The wage function to be estimated is given by:

$$\ln y_i = \beta_0 + \beta_1 x_{1i} + \beta_2 q_i + \beta_3 r_i + \varepsilon_i + \phi_{1i} q_i + \phi_{2i} r_i$$

where $x_1$ represents other determinants of earning, and $\phi_1$ and $\phi_2$ are unobservable individual heterogeneity effects that depend on fatal and non-fatal risks. For instance, individual attributes such as cool-headedness may make an individual more productive, or less at risk, in a riskier situation. As he puts it, "Because *q* and *r* are choice variables of individuals, the chosen value of *q* can depend on, and thus be correlated with, $\varepsilon$, $\phi_1$ and $\phi_2$. Therefore, OLS may provide biased estimates." (p. 10) The suggested solution is to estimate the equation jointly with regression equations for mortality risk and morbidity risk. These two equations are:

$$q = \pi_0 + \pi_1 x_1 + \pi_2 x_2 + \pi_3 z + \eta$$

$$r = \delta_0 + \delta_1 x_1 + \delta_2 x_2 + \delta_3 z + \mu$$

where $x_2$ is a variable which proxies the degree of risk aversion and $z$ is non-wage income.

The identification problem discussed by Garen (1988) is intrinsic to the endogenous risk setting, and plagues much of the empirical literature on risks in a natural setting. An attractive alternative is the use of controlled laboratory experiments to tease out the effect of endogeneity on behavior. Very few laboratory experiments, however, have been designed to study the endogeneity issue.[8] The striking exceptions are Shogren (1990) and Shogren and Crocker (1994)(SC). The basic model in each case is a simple lottery with probability *p* of loss *L* and *(1-p)* probability of gain *G*. The EU of this lottery is:

$$\text{EU}= pU(M-L)+(1-p)U(M+G)$$

In their experiments self-protection is the certainty equivalent of this lottery, which is:

$$U(M+G-x) = pU(M-L)+(1-p)U(M+G)$$

This is the amount an individual is willing to pay with certainty to completely reduce the loss to zero. Self-protection expenditure in this case guarantees a positive gain.

Self-insurance in these experiments is the payment made in all states of nature to reduce loss to zero. That is, self-insurance is given by:

$$pU(M-z)+(1-p)U(M+G-z) = pU(M-L)+(1-p)U(M+G)$$

Notice that although self-protection expenditure guarantees a gain, self-insurance merely reduces loss to zero. It seems logical that self-protection will be preferred to self-insurance in

[8] There have been many experiments that examine risk mitigation choices, but not always in the context in which the risk mitigation affects the risk of the outcomes. For example, in the context of forest fire risks, McKee, Berrens, Jones, Helton and Talberth (2004) and Talberth, Berrens, McKee and Jones (2006) show that people are willing to pay positive amounts of money for averting wildfire, and insuring against losses from wildfire, even in the presence of market insurance. They employ a combination of experimental and contingent valuation data to establish the existence of self-insurance investment.

this setup. In our design self-protection does not guarantee a gain: it only reduces the probability of a loss. This makes it more comparable to self-insurance.

The primary objective of the SC design is to test whether use of alternative risk-reduction mechanisms lead to difference in choices. In their design they have four different mechanisms available for mitigating risk: individual self-protection, individual self-insurance, collective self-protection and collective self-insurance. They find the maximum willingness to pay for individual self-protection and the minimum willingness to pay for collective self-insurance. SC recognize that the difference in these two valuations may be partly due to the presence of a free rider problem in the collective data, because of the public goods nature of collective risk. SC finds that risk-reduction mechanisms matter, and conclude that there does exist a framing effect on risk mitigation behavior.

In SC a Vickrey sealed bid second price auction is used for the valuation of private risk reduction mechanisms. Each subject competes for self-protection or insurance. The highest bidder gets the good at the price of the second highest bid. A "Smith auction" is used for the collective risk reduction mechanisms.[9] Each subject makes a bid to reduce risk, and if the sum of the bids exceeds or equals the price, the good is provided. The use of market mechanisms may result in strategic bidding, in which the true value of risk-reduction may not be revealed even when truthful revelation is a dominant strategy (Rutström (1998) and Harstad (2000)). Simple individual choice mechanisms, like choice between lotteries, are simpler to explain and have no potential for strategic bidding. This is why we use individual lottery experiments.

[9] In a Smith auction (1980) each individual has to accept his share of cost by bidding that amount. The public good is only provided when there is unanimity about the quantity to be provided.

## II. Experimental Design

The experimental design is deliberately structured to allow conceptual and econometric identification of the effects of endogenous risk. The subject faces 11 decision tasks: a standard lottery task in the gain domain, a standard lottery task in the loss domain, four lottery tasks with endogenous probabilities in the loss domain, three WTP tasks, and two betting tasks. The subject picks one task at random out of the 11 that determines his earnings.[10]

Each of the 11 tasks involves a series of binary choices. Once the task for payment has been decided, the subject is paid for one of these binary choices selected at random by rolling a die. This has the advantage of avoiding any "wealth or portfolio effect" that may arise if subjects are paid for multiple decisions made at the same time.[11]

The betting tasks and the WTP tasks are undertaken as part of a VR experiment. This VR experiment studies the subject's perception of the risk of forest fire in a betting task, and their WTP to reduce this risk. The risk of forest fire is carefully explained to the subject using different methods, including an opportunity to experience computer simulations and graphical rendering of forest fires.

---

[10] A simple procedure is used to ensure that the subject believes that the process of choosing this task is random, using the following instructions: "The box in front of you has 11 envelopes and 11 cards numbered 1 through 11. Please put one card into each envelope and close the envelopes. I will now shuffle these envelopes. The envelopes have now been carefully shuffled and I ask that you pick one of them. The number on the card in the envelope you selected determines which task you will be paid for. But you will not know which one until the end of the experiment when you will be allowed to open the envelope." All experiments were conducted with one subject at a time.

[11] The experimental procedure of picking only one choice for payment, to avoid portfolio effects, has been widely used. It is not theoretically obvious that portfolio effects should always be avoided when eliciting risk attitudes, as explained in Harrison and Rutstrom (2008, p.117). For our purpose however, it is desirable to avoid any portfolio effects since payment for the betting tasks as well as choice tasks in VR are determined by the outcome of simulations of a forest fire with identical underlying distributions.

A large part of our instrument deals with the instructions for the VR experiment, including the betting mechanism. Subjects first face the VR tasks, followed by the lottery tasks. There are two alternative treatments of the VR experiment, to check for order effects: either the subject faces the betting task first followed by the WTP task or the subjects face the WTP task first followed by the betting task.

First, we describe the VR scenario and the WTP table that the subject was given. Second, we describe the betting mechanism we use to directly elicit beliefs from subjects. This is followed by a description of the various lottery tasks we use in our experiment: the standard Holt and Laury (2002) task in the gain domain, followed by similar tasks in the loss domain, and then by lottery tasks with endogenous risk.

### A. WTP task in the VR experiment

The design of the VR choice task is taken from FHHR. Following their design, we pay subjects according to damages to their personal property, which is simulated as a virtual log cabin at the edge of Ashley National Forest, Utah. The subject is given an initial cash credit. With this credit he has the option of paying for a prescribed burn policy which would reduce the risk of his property burning down in the event of a fire. But this policy costs him some money up front – non-refundable in the event that his property does not burn. FHHR view fire management options as policy lotteries:

> Forest fire management options such as prescribed burn can be viewed as choices with uncertain outcomes. In the language of economists these options are "lotteries" or "prospects" that represent a range of final outcomes, each with some probability, and hence can be viewed as a policy lottery. One outcome in this instance might be "my cabin will not burn down in the next 5 years" or "my cabin will burn down in the next 5 years." Many more outcomes are obviously possible: these are just examples to illustrate concepts. (p. 73)

Using their experimental design, the choice in the policy lotteries presented to subjects is between a "risky lottery" with no prescribed burn and a "safe lottery" with a prescribed burn. If the risky lottery is chosen, and therefore no prescribed burn is undertaken, the probability of the house being destroyed by wild fire is relatively high. Undertaking a prescribed burn, or choosing the safe lottery, reduces the probability of the house being destroyed by fire. In short, the subject chooses the level of risk that he is willing to take, where risk is endogenous and chosen by the subject rather than being exogenously given to him.

Choosing a lower risk lottery, however, involves an upfront cost of prescribed burning. Prizes in both states of nature, where the house burns down and the house does not burn down, are higher for the risky lottery than for the safe lottery. But the probability of the worst outcome is higher in the risky lottery, implying that there exists a tradeoff between the level of risk that a subject faces and the amount of money he is willing to pay to reduce that risk. This is typical of models with endogenous risk where an individual is able to affect the level of risk he faces by investing in mitigation.

In this VR, the subject is able to experience computer simulations of the forest fire in order to form his own beliefs about the probability of damage. As explained in FHHR:

> Contrary to decisions involving actual lottery tickets, a policy lottery has outcomes and probability distributions that are not completely known to the agent. We therefore expect the choices to be affected by individual differences not just in risk preferences, but also in risk perceptions. (p. 73)

Using a separate betting mechanism, described below, we directly elicit this risk perception.

We are interested in the perception of risk that subjects form in this VR experiment. The instructions are designed to convey information about this risk as accurately as possible. First, we give a brief introduction to the subject about the threats of wild fires. We then introduce prescribed burning as a fire management tool that can reduce the frequency and severity of fires. The idea of VR computer simulation is introduced, with explanations that the predicted path and severity of the fire depend on things such as topography, weather, vegetation, and ignition points. Next we give the instructions for the first experimental task, which is alternatively the choice task or the betting task. This is followed by information about the risks of wild fires. Subjects are made familiar with the idea that fires and fire damages are stochastic, and can be described through frequency distributions. The distributions that are presented to them are generated through Monte Carlo simulations using a computer simulation model, FARSITE, used by the U.S. Forest Service for this very purpose. The subjects then experience 4 dynamic VR simulations of specific wild fires, 2 for each of the cases with and without previous prescribed burns, rendered from the information supplied by FARSITE simulations that vary weather and fuel conditions. We selected these simulations to represent a high and low risk of fire damage, and the subjects are told this.

Subjects are told that the VR simulations are based on the Ashley National Forest in Utah, and that they have a virtual property in this area. The simulated area is subject to wild fire, and they must make a decision whether to pay for a prescribed burn or not, which would reduce the risk that their property would burn. The information about risks to their property that subjects receive is threefold.

First, they are told that the background uncertainties are generated by (a) temperature and humidity, (b) fuel moisture, (c) wind speed, (d) duration of the fire, and (c) the location of the ignition point. They are also told that these uncertainties are binary for all but the last, which is ternary; hence there are 48 background scenarios. They are also told the specific values for these conditions that are employed (e.g., low wind speed is 1 mph, and high wind speed is 5 mph). The instructions explain that these background factors will affect the wild fire using a computer simulation model developed by the U.S. Forest Service to predict the spread of wild fire. Thus subjects could use this information, their own sense of how these factors play into wild fire severity, and their experiences and inferences from the VR experience (explained below), to form some probability judgments about the risk to their property. The objective is to provide information in a natural manner, akin to what would be experienced in the actual policy-relevant choice, even if that information does not directly "tell" the subject the probabilities.

Second, the subject is shown some histograms displaying the distribution of acreage in Ashley National Forest that is burnt across the 48 scenarios. Figure 1 shows the histograms presented to subjects. The scaling on the vertical axis is deliberately in terms of natural frequencies defined over the 48 possible outcomes, and the scaling of the axes of the two histograms is identical to aid comparability. The qualitative effect of the enhanced prescribed burn policy is clear: to reduce the risk of severe wild fires. Of course, the information here is about the risk of the entire area burning, and not the risk of their personal property burning, and that is pointed out in the instructions.

Third, they are allowed to experience several of these scenarios in a VR environment that renders the landscape and fire as naturally as possible. Figure 2 illustrates the type of graphical rendering provided, although static images such as these do not do justice to the VR "presence" that was provided. Some initial training in navigating in the environment is provided, which for this software is essentially the same as in any so-called "first person shooter" video game.[12] The mouse is used to change perspective and certain keys are designated for forward, backward, and sideward movements. The subject is then presented

[12] For student subjects this interface is second nature, and occasionally their Second Life.

with the 4 scenarios and is then free to explore the environment, the path of the fire, and the fate of their property during each of these. Apart from the ability to move across space, subjects also have the option of moving back or forth in time within each fire scenario.[13]

Table 1 shows the comparable payoff matrix that is implied by our VR instrument, assuming that the subject accurately infers the true probabilities of his own property burning. From the table we see that the true probabilities of the cabin burning are 0.06 if the prescribed burn policy is chosen and 0.29 if the subject chooses not to pay for the policy. All prizes are stated as losses from an initial credit, here $40. Lottery B gives a 29% chance of a loss of $18 and a 71% probability of no loss. Lottery B remains the same in all rows of this table, with a fairly high probability (0.29) of the property being damaged in a wild fire. It costs nothing if no damage occurs,[14] and this is the risky lottery. Lottery A is the safe lottery: in the first row, it offers only a 6% probability of a loss of $18, while there is a 94% probability of no loss occurring. Clearly, lottery A is preferred to lottery B in the first row. In the second row, lottery A has a 6% probability of a loss of $20 and a 94% probability of a loss of $2. In this case it costs $2 to undertake prescribed burning that reduces the probability of damage from 29% to 6%, but also results in both prizes being lower than in the risky lottery. As one moves down the table, the amount of money paid for prescribed burn increases and accordingly both prizes in lottery A decrease. Following this logic, the last row of this table corresponds to the maximum cost of prescribed burning, which, in our experiment, equals the value of the house. In this example the cost of prescribed burn varies

[13] This points to another feature of the VR environment in settings where current action, or inaction, can lead to latent effects well into the future. The VR simulation interface can be used by subjects to "fast forward" and better comprehend those effects. The credibility of the future-generated scenarios is conditional on the credibility of the modeling of dynamic processes in the underlying simulation model, of course.

[14] The probabilities 0.29, 0.71, 0.06 and 0.94 are the ones used in the computer simulations of the wild fires.

from $0 to $18 (the value of the simulated house). This makes the instrument intuitive for subjects to follow and ensures that the switch point, from safe to risky lottery, falls within the range of the table.

For each choice task we present the subject with an array of possible costs for the policy, between $0 and $18 in increments of $2. In each case the subject is asked to simply say "yes" or "no" as to whether they wanted to contribute to the policy at that cost.[15] The subject is told that one of these costs would be selected at random, using a 20-sided dice, and if the prescribed burn had been selected for that cost it would be put in place, otherwise not. After the actual cost was selected at random, the subject was told that we would then select each of the background factors by rolling a die, so that one of the 48 background scenarios would be selected with equal probability. Thus the subject's choices, plus the uncertainty over the background factors affecting the severity of a fire conditional on the chosen policy, would determine final earnings.

There are 3 such decision tasks that the subject faces. One is when the damage to the house results in a loss of $38 with an initial credit of $80, one is where the damage is $28 with an initial credit of $60, and the final one is where the damage is $8 with an initial credit of $20. The choice data generated in these three tasks can be analyzed as data from a Multiple Price List of discrete pairwise lottery options, apart from the fact that the subjective probabilities are not known and must be estimated. Using the betting task, described below, and this choice task we can jointly estimate subjective perception of risk and risk attitude.

[15] Instructions and decision sheets are available on request.

### B. Betting task

The objective of the betting task is to directly recover the subject's belief that event A will occur instead of event B, where A and B are mutually exclusive. Assume that the subject is risk neutral and has no stake in whether A or B occurs other than the bets being made on the event. Let the subject be told that there are 9 bookies, each willing to take a bet at stated odds. Table 2 shows odds for the two events in the *form* that they are naturally stated in the field: what is the amount that the subject would get for a $1 bet if the indicated event occurred?

In our design the subject is simply asked to decide how they want to bet with each of the 9 bookies: do they want to bet on A or B? Their "switch point," over the 9 bookies, is then used to infer their subjective belief. The basic experimental design and estimation strategy are borrowed from FHHR and Andersen, Fountain, Harrison and Rutström (2009a).[16] Consider a subject that has a personal belief that A will occur with probability 0.75, and assume that the subject *has* to place a bet with each bookie, knowing that only one of these bookies will actually be played out. The odds offered by a particular bookie are shown on a given row, so different rows correspond to different bookies. The (risk neutral) subject would rationally bet on A for every bookie offering odds that corresponded to a lower house probability than 0.75 of A winning, and then switch over to bet on B for every bookie offering odds that corresponded to a higher house probability than 0.75 of A winning. These bets are shown in Table 2, and imply gross earnings of $10 or $0 with the

[16] Familiar scoring rule procedures are formally identical, since each probability report implicitly generates a bookie willing to bet at certain odds. Thus when the subject makes a report in a Quadratic Scoring Rule, for example, the subject is in effect choosing to place a bet on the event occurring with payoffs given by odds that are defined by the scoring rule. By making one report instead of another, the subject is then choosing one bet over another, or equivalently in our design, one bookie over another.

first bookie, $5.00 or $0 with the second bookie, and so on. The expected gross earnings from each bookie can then be calculated using the subjective belief of 0.75 that the subject stated. Hence the expected gross earnings from the first bookie are (0.75 × $10) + (0.75 × $0) = $7.50, and so on for the other bookies. A risk neutral subject would bet on event A for the first 7 bookies and then switch to event B.

Each subject faces 2 betting tasks in our experiment. The first event is that the house in the VR environment burns down when the fuel load is low (or a prescribed burn is used); and the second event is that the house burns down when fuel level is high (or no prescribed burn is used). The subject is given a $5 stake to bet with in each of the 2 betting tasks, and a bet has to be placed for each of the 9 bookies. The $5 from one house is not transferable to other houses, and one of the bets will be selected at random to be actually played out. In effect, we "force feed" the subjects by requiring that they place a bet with each betting house, and do not allow them to change the $5 stakes: all the subject can do is decide if they want to bet on event A or event B in each house.

### C. Standard Lottery Task in Gains Domain

Holt and Laury (2002) devise a simple experimental measure for risk aversion using a Multiple Price List design. Each subject is presented with a choice between two lotteries, which we can call lotteries A or B. Table 3 illustrates the basic payoff matrix presented to subjects in our experiment.[17] The first row shows a choice between getting $24 for certain and $1 for certain. The second row shows that lottery A offers a 90% chance of receiving $24 and a 10% chance of receiving $26. The expected value of this lottery is shown as

[17] The prizes for this standard lottery task have been adjusted to match our other tasks. Here the prizes are $1, $24, $26 and $50 instead of the original $3.85, $2.00, $1.60 and $0.10 in Holt and Laury (2002)

$24.20, although the EV columns are not presented to subjects. Similarly, lottery B in the second row has the prizes $50 and $1, for an expected value of $5.90. Thus the two lotteries have a difference in expected values of $18.30. As one proceeds down the matrix, the expected value of both lotteries increase, but the expected value of lottery B becomes greater relative to the expected value of lottery A.

Each subject chooses A or B in each row, and one row is later selected at random for payment. The logic behind this test for risk aversion is that only risk loving subjects would take lottery B in the second row, and only very risk-averse subjects would take lottery A in the last row. Arguably, the first row is simply a test that the subject understood the instructions, and has no relevance for risk aversion at all. A risk neutral subject should switch from choosing A to B when the EV of each is about the same, so a risk-neutral subject would choose A for the first five rows and B thereafter. The standard lottery task is an example of exogenous risk. In each row, the subject is equally likely to get the bigger prize or the smaller prize. The only difference is the prize amount.

We take each of the binary choices of the subject as data, and estimate the parameters of a latent utility function that explains those choices using an appropriate error structure to account for the panel nature of the data. Once the utility function is defined, for a candidate value of the parameters of that function, we can construct the expected utility of the two gambles, and then use a linking function to infer the likelihood of the observed choice. Extensions to non-EUT specifications are immediate.

### D. Standard Lottery Task in the Loss Domain

The standard lottery task in the loss domain presents to the subject the same lottery choices discussed in the previous section, except that now they are framed as losses instead of gains. The basic hypothesis to be tested is that the risk aversion coefficient[18] in the gain and loss domains are identical. Holt and Laury (2008) report the results of similar experiments where prizes are stated as losses instead of gains. In their experiments the prizes in the loss domain are "reflections" of prizes in the gain domain, which simply means that the prize amounts remain unchanged but the signs are reversed. That is, instead of winning $16 the subject now loses $16 with the same probability as before. For the example considered in Table 3, the prizes after reflections would become -$24, -$26, -$0 and -$49.[19]

Our primary interest in restating the standard lottery task in the loss domain is to make it comparable with other tasks in our experiment. The VR task, discussed earlier, is different from the standard lottery task because the former deals with a loss from an initial credit, while the latter is designed as a gain with no initial credit. In the lottery task in the loss domain, the subject is given an initial credit of $50. Any loss is deducted from this initial credit, and the remaining amount is the subject's net earnings. For example, in the first row of Table 4 Lottery A offers a loss of $26 with certainty and lottery B a loss of $49 with certainty. The second row offers a choice between lottery A with a 90% chance of losing $26 and a 10% of losing $24 and lottery B with a 90% chance of losing $49 and a 10% chance of losing $0. To a risk neutral subject, with a reference point of $0, this is the same lottery choice as the one given in the first row of Table 3, since lottery A gives a 90% chance of getting $50 – $26 = $24 and a 10% chance of getting $50 - $24 = $26; and lottery B gives a 90% chance of getting $50 - $49 = $1 and a 10% of getting $50 - $0 = $50.

[18] Or coefficients, in the case of EUT or non-EUT specifications that require the estimation of two or more structural parameters to characterize risk attitudes.
[19] Holt and Laury (2008) had their subjects earn their initial endowment in earlier experimental tasks, which averaged $43 and ranged from $21.68 to $92.08.

These lotteries are, however, identical only to the extent that the assumption about the reference point is valid. As emphasized in Harrison and Rutström (2008; p.95ff.), it is very difficult to determine the right reference point that the subject actually uses in a risk aversion task. It is possible, for instance, that the subject has a reference point of $50, and integrates the $50 credit into his wealth stream. In this case he views lottery A as a loss of $26 rather than a gain of $24. The difference in expected value of the lotteries for both $0 and $50 reference points are identical. Hence in both cases a risk neutral subject chooses lottery A for the first five rows and lottery B for the last five rows.[20]

Arguably, the reference point does not matter when comparing tasks in the loss domain where the initial credit and show up fee are identical, as long as reference points are not task specific. When comparing lotteries between the gain and loss domain, however, the indeterminacy of the reference point remains central. If the reference point is $0, then the standard lottery choices in the gain domain and loss domain are identical and we expect behavior to be identical. If, on the other hand, the reference point is $50 then the prize in Table 4, for instance, is viewed as a loss of $26 and not as a gain of $24. Allowing for loss aversion, it is possible that the subject makes different choices in Table 3 and 4. We test whether estimated risk aversion coefficients in the gain and loss domains are identical and whether there is evidence of framing effect.

[20] Of course, $0 and $50 are not the only two possible reference points. If the subject integrates the show up fee of $5 into his wealth coefficient, then the reference points are $5 ands $55 respectively. If he integrates his lifetime income or income outside the current experiment things get more complicated. Inferences about lotteries in the loss domain are very sensitive to assumptions about the reference point.

### E. Lottery Task in the Loss Domain with Endogenous Risk

The lottery task with endogenous risk is the lottery version of the VR choice task. Table 1 describes the policy lottery that the subject faces, assuming that the subject is able to correctly infer the risk of the house burning down in the simulated forest fire. In the current task, the policy lottery is presented to the subject simply as a lottery choice without any of the contexts and visuals of the VR choice task. Notice though, that in this lottery version subjects are given the probabilities of the different prizes, while in the VR experiments they have to form their own beliefs.

Each subject faces 4 lottery tasks with endogenous risk in the loss domain. The maximum WTP varies with the initial credit given. When the initial credit is \$80 (or \$20), the maximum WTP or value of the house is \$38 (or \$8).[21] The probability of incurring a loss in Lottery B, the risky lottery, is alternatively 0.29 or 0.59; the probabilities for the safe lottery remain unchanged.

Having both the standard lottery task and the lottery task with endogenous risk allows us to test for framing effect when risk is exogenous and endogenous. We test whether estimated risk aversion coefficients and subjective beliefs, recovered from the two tasks, are identical.

The sample consists of 35 subjects recruited from the student population of the University of Central Florida. All experiments were conducted in May and June, 2009.

---

[21] In the WTP task in the VR experiment we have 3 treatments: initial credit of \$80 with house valued at \$38; initial credit of \$60 with the house valued at \$28; initial credit of \$20 with the house valued at \$18. In the lottery task with endogenous risk, we do not include the treatment with an initial credit of \$60 and maximum WTP of \$28.

## III. Results

We have two kinds of tasks – lottery tasks and VR WTP tasks. From the lottery tasks we get an estimate of risk attitude, based on the concavity of the utility function in the EUT model. In the VR tasks the subject is able to experience computer simulations of forest fires in order to form his own beliefs about the probability of damage to a property in the event of a forest fire. Therefore, we are able to recover the subject's beliefs as well as risk attitude from these tasks.

We first review the method of estimating risk attitude using EUT. Then we discuss the procedure for jointly estimating belief and risk attitude using the WTP and the betting mechanisms, maintaining the assumptions of SEUT.

### A. Estimating Risk Attitude under EUT

Assume for the moment that utility of income is defined by

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

where x is the lottery prize and $r \neq 1$ is a parameter to be estimated.[22] Thus, *r* is the coefficient of constant relative risk aversion (CRRA): $r = 0$ corresponds to risk neutrality, $r < 0$ to risk loving, and $r > 0$ to risk aversion. Let there be *K* possible outcomes in a lottery. If the probability of outcome $k$ is $p_k$, expected utility (EU) is simply the probability weighted utility of each outcome in each lottery *i*:

$$EU_i = \sum_{k=1,K} (p_k \times U_k)$$

The choice depends on the difference in EU between the right and the left lotteries.

$$\nabla EU = EU_R - EU_L$$

[22] For r=1, assume U(x)=ln(x) if needed.

where $EU_R$ is the "right" lottery and $EU_L$ is the "left" lottery. This latent index, based on latent preferences, is then linked to the observed choices using a standard cumulative normal distribution function $\Phi(\nabla EU)$. This "probit" function takes any argument between $\pm\infty$ and transforms it into a number between 0 and 1 using the function shown in Figure 3. Thus we have the probit link function,

prob (choose lottery R) = $\Phi(\nabla EU)$

This function forms the link between the observed binary choices, the latent structure generating the index, and the probability of that index being observed. The logistic function is very similar to the probit, as illustrated in Figure 3, and leads instead to the "logit" specification.

In the standard lottery task the latent index is the difference in EU between the right and left lottery, and the observed choice is the choice of the right or left lottery. The conditional log-likelihood function is:

$$\ln L(r; y.X) = \sum_i [(\ln \Phi(\nabla EU) \times I(y_i = 1)) + (\ln(1 - \Phi(\nabla EU)) \times I(y_i = -1))]$$

where $I(.)$ is the indicator function and $y_i = 1$ (-1) indicates the choice of right (left) lottery. The only variable that has to be estimated from this log-likelihood function is *r*.

An important extension of the core model is to allow subjects to make errors in the decision process. The notion of error is one that has already been encountered in the form of the statistical assumption that the probability of choosing a lottery is not 1 when the EU of that lottery exceeds the EU of the other lottery. This is implicitly assumed when one adopts a link function, of the kind shown in Figure 3, to go from the latent index to observed choices. The contextual error specification, suggested by Wilcox (2009), introduces a normalizing term $\nu$ for each lottery pair, and a structural "noise parameter" $\mu$ to allow for error from the deterministic EU model:

$$\nabla EU = [(EU_R - EU_L)/\nu]/\mu$$

The normalizing term $\nu$ is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair, and ensures that the *normalized* EU difference $(EU_R - EU_L)/\nu$ remains in the unit interval. This normalization allows one to define robust measures of "stochastic risk aversion," in parallel to the deterministic concepts from traditional theory. Notice that when $\mu = 1$ we return to the original specification without error. As $\mu$ increases, the above index falls until, at $\mu = \infty$, it collapses to zero, so that the probability of either choice becomes ½. In other words as the noise in the data increases, the model has less and less predictive power until at the extreme the prediction collapses to fifty-fifty or equal likelihood of both choices.

To allow for subject heterogeneity with respect to risk attitudes, the parameter $r$ is modeled as a linear function of observed individual characteristics of the subject. For example, assume that we only had information on the age and sex of the subject, denoted **age** (in years) and **female** (0 for males, and 1 for females). Then we would estimate the coefficients $\alpha$, $\beta$ and $\eta$ in $r = \alpha + \beta \times \text{age} + \eta \times \text{female}$. The covariates we use are all binary, except one, and reasonably self-explanatory. The variable **wtp** is a dummy for whether (=1) or not the task involves a willingness to pay table, and hence reflects an endogenous risk setting; **loss** is a dummy variable for whether (=1) or not the task is presented in the loss frame; **betfirst** is a dummy for whether the WTP tasks came first (=1) or not; **age** is given in years over 17; **female** is a dummy for whether the subject is female (=1) or male; **hispanic** is a dummy for Hispanic heritage (=1); **business** is whether or not the subject is a business major (=1); **GPAhigh** is for subjects with a self-reported GPA higher than 3.24 (=1); **works** is a dummy for whether (=1) or not the subject is employed.

In this design multiple responses are elicited from each subject. This may lead to clustering, or heteroskedasticity. Therefore while estimating the model it is essential to correct for clustering effects.

Using this specification on the lottery tasks, we find no evidence of framing effects, and evidence of modest risk aversion, consistent with a large body of existing literature (e.g., Holt and Laury (2002)(2005), Harrison, Johnson, McInnes and Rutström (2005) and Harrison and Rutström(2008)). Maximum likelihood estimates of the EUT model are presented in Table 5.[23] There is no framing effect of the task being presented in the loss frame rather than the familiar gain frame of standard lottery tasks.[24] Nor is there any evidence of an order effect between lottery tasks and betting tasks. The average predicted CRRA value for this sample is 0.32, implying modest risk aversion, but there is sufficient heterogeneity in the sample that one cannot rule out risk neutrality on average (the 95% confidence interval on the predicted CRRA is between -0.37 and 1.01). It would be inappropriate to assume risk neutrality for all subjects, because we want the risk attitudes of different subjects to condition their inferred subjective beliefs; to assume $r = 0$ for every subject, with *no standard error*, would mis-characterize the true distribution of risk attitudes across the sample.

[23] We also evaluated extensions of the CRRA utility specification to allow for increasing or decreasing RRA over the prize domain. In this case we used the expo-power specification employed by Holt and Laury (2002). We find no evidence of non-CRRA behavior in this domain. Nor are there any differences in the qualitative conclusions about framing effects and the role of endogenous risk. For these reasons we use the simpler CRRA specification throughout.

[24] One could extend this analysis to include a formal estimation of loss aversion, but for reasons explained at length in Harrison and Rutström (2008; §3.2.3) and Harrison and Rutström (2009), we do not want to do so casually.

**B. Jointly Estimating Belief and Risk Attitudes under SEUT**

The responses to the belief elicitation task can be used to estimate the subjective probability that each subject holds if we are willing to assume something about how they make decisions under risk.

If the subject is assumed to be risk neutral, then we can directly infer the subjective probability from the report of the subject. When allowing for varying risk attitudes, we jointly estimate the subjective probability and the parameters of the core model. The subject that selects report $\theta$ receives the EU given by

$$EU_\theta = \pi_A \times U \text{ (payout if A occurs } | \text{ report } \theta)$$

$$+ (1-\pi_A) \times U \text{ (payout if B occurs } | \text{report } \theta)$$

where $\pi_A$ is the subjective probability that A will occur. The example in the second row of Table 1 now becomes

$$EU_{safe} = (1-\pi_{safe}) \times (38)^{(1-r)} /(1-r) + \pi_{safe} \times U(20)^{(1-r)} /(1-r)$$

and

$$EU_{risky} = (1-\pi_{risky}) \times (40)^{(1-r)} /(1-r) + \pi_{risky} \times (22)^{(1-r)} /(1-r)$$

where $\pi_{safe}$ is the subjective probability of the house burning down when prescribed burning is implemented and $\pi_{risky}$ is the subjective probability that the house will burn down if no prescribed burning is implemented. More generally,

$$EU_{safe} = \pi_{safe} \times U\text{(payout if house burns } | \text{ pay for prescribed burn)}$$

$$+ (1-\pi_{safe}) \times U\text{(payout if house does not burn } |\text{pay for prescribed burn)}$$

and

$$EU_{risky} = \pi_{risky} \times U(\text{payout if house burns} \mid \text{do not pay for prescribed burn})$$

$$+ (1 - \pi_{risky}) \times U(\text{payout if house does not burn} \mid \text{do not pay for prescribed burn}).$$

The latent index in this problem is the difference in EU from paying for prescribed burning and not paying for prescribed burning:

$$\nabla EU = [(EU_{safe} - EU_{risky}) / \nu] / \mu$$

Apart from $r, \nu$ and $\gamma$ we now need to estimate the two beliefs $\pi_{safe}$ and $\pi_{risky}$. The *joint* maximum likelihood problem is to find the values of these parameters that best explain observed choices in the belief elicitation tasks as well as observed choices in the lottery tasks.

Detailed estimates are contained in Tables 6 and 7, and displayed in Figures 4 and 5. In addition to the same covariates as in Table 5, we also employ interactions of the central **wtp** dummy with demographic characteristics. The estimates of the subjective probabilities reported here as pLO and pHI refer to a standard non-linear, log-odds transformation[25] to ensure that the latent probabilities $\pi_{safe}$ and $\pi_{risky}$, respectively, fall between 0 and 1 while the estimated parameters pLO and pHI are unbounded; hence a value of 0 reflects a probability of ½, and a positive (negative) coefficient reflects a probability below (above) ½.

Table 6 assumes that the same risk attitudes and subjective beliefs apply to the exogenous and endogenous risk settings, and Table 7 allows each to vary with the setting. Our main hypothesis is that there is no effect from allowing these to vary with the setting. Figures 4 and 5 display kernel density functions of the predicted subjective probability for each subject in each setting. These densities do not reflect the standard errors in the

[25] If we estimate the parameter $\tau$ then the transformation to the probability $\pi$ is $\pi = 1/(1+\exp(\tau))$. So $\pi_{safe} = 1/(1+\exp(\text{pLO}))$ and $\pi_{risky} = 1/(1+\exp(\text{pHI}))$.

parameter values, which are shown in Tables 6 and 7. Hence formal hypothesis tests must rely on the quantitative estimates, and not these pictures.

We find evidence of approximately linear utility functions, although the variability across the sample is large. We do observe a smaller CRRA coefficient when there is endogenous risk, as shown by the coefficient on the dummy variable **wtp** in Table 7, but this is not statistically significantly different from zero and has a *p*-value of 0.68. So we conclude that risk attitudes are not source dependent.

Turning to subjective beliefs, we find a statistically significant effect from endogenous risk on each of the two subjective probability values estimated. Formally these involve hypothesis tests of the joint restriction that all of the **wtp**-interaction terms on pHI or pLO are jointly and simultaneously zero. The null hypothesis that these coefficients are all zero can be rejected for pHI, for pLO, and for pHI and pLO taken together. Figure 5 shows that the overall effect is to make the subjective belief distributions move towards the extremes as we allow for endogenous risk settings. It is as if the option of risk mitigation is *subjectively* moving a risky environment towards a non-risky environment, which is exactly what one wants to do when "managing" risk. On the other hand, the conversion to a non-risky environment is not complete, nor should it be if it is too subjectively costly in terms of EU to do so.

## IV. Conclusion

Risk attitudes and subjective beliefs are two fundamental determinants of decision making under risk. Virtually all settings in which there is risk involve what we refer to here as endogenous risk, rather than the exogenous risk that is commonly studied. Our results provide evidence that behavior is different in these two environments. There do not appear to be significant differences in attitudes to risk, but subjective beliefs are very different. Our design provides the simplest setting in which these structural effects can be untangled.

One minor extension of our approach would be to move from the SEUT framework to one in which subjective probabilities are weighted by some increasing function, and decision weights employed on the utility of outcomes, following Kahneman and Tversky (1979) and Quiggin (1992).[26] This extension would be valuable because what appear to be changes in subjective probabilities assuming SEUT, and hence assuming no such probability weighting, might in fact be a change in the probability weighting function for unchanged subjective probabilities.

One major extension would be to relax the maintained assumption that agents behave as if their subjective beliefs are statistically degenerate in the sense of being probabilities that are held with certainty. If they are, instead and plausibly, uncertain beliefs, then we need to account for uncertainty aversion.[27] It is possible that what appear to be shifts in subjective probabilities, conditional on unchanged risk attitudes, are confounded by not modeling attitudes towards uncertain subjective beliefs. In any event, our results show that one cannot simply assume that risk attitudes and beliefs elicited in an exogenous risk setting transfer to endogenous risk settings.

Our design and analysis provides the basis for further exploration of the structural determinants of the difference in behavior towards exogenous and endogenous risk.

---

[26] Harrison and Rutström (2008, §3.1) show how our approach can be used to generate estimates of risk attitudes that consider probability weighting, and Andersen, Fountain, Harrison and Rutström (2009a) show how this extends immediately to the estimation of subjective probabilities from a belief elicitation task.

[27] Andersen, Fountain, Harrison and Rutström (2009b) explore the theoretical, experimental and econometric issues involved here.

**Table 1. Inferred WTP Instrument when Initial Credit is $40**

**and the House is Valued at $ 18**

| Lottery A | | | | Lottery B | | | | $EV^A$ | $EV^B$ | Difference |
|---|---|---|---|---|---|---|---|---|---|---|
| p(safe) | | p(burn) | | p(safe) | | p(burn) | | | | |
| 0.06 | -18 | 0.94 | 0 | 0.29 | -18 | 0.71 | 0 | -1.08 | -5.22 | 4.14 |
| 0.06 | -20 | 0.94 | -2 | 0.29 | -18 | 0.71 | 0 | -3.08 | -5.22 | 2.14 |
| 0.06 | -22 | 0.94 | -4 | 0.29 | -18 | 0.71 | 0 | -5.08 | -5.22 | 0.14 |
| 0.06 | -24 | 0.94 | -6 | 0.29 | -18 | 0.71 | 0 | -7.08 | -5.22 | -1.86 |
| 0.06 | -26 | 0.94 | -8 | 0.29 | -18 | 0.71 | 0 | -9.08 | -5.22 | -3.86 |
| 0.06 | -28 | 0.94 | -10 | 0.29 | -18 | 0.71 | 0 | -11.08 | -5.22 | -5.86 |
| 0.06 | -30 | 0.94 | -12 | 0.29 | -18 | 0.71 | 0 | -13.08 | -5.22 | -7.86 |
| 0.06 | -32 | 0.94 | -14 | 0.29 | -18 | 0.71 | 0 | -15.08 | -5.22 | -9.86 |
| 0.06 | -34 | 0.94 | -16 | 0.29 | -18 | 0.71 | 0 | -17.08 | -5.22 | -11.86 |
| 0.06 | -36 | 0.94 | -18 | 0.29 | -18 | 0.71 | 0 | -19.08 | -5.22 | -13.86 |

**Figure 1. Histogram Displaying Distribution of Forest that Burned**

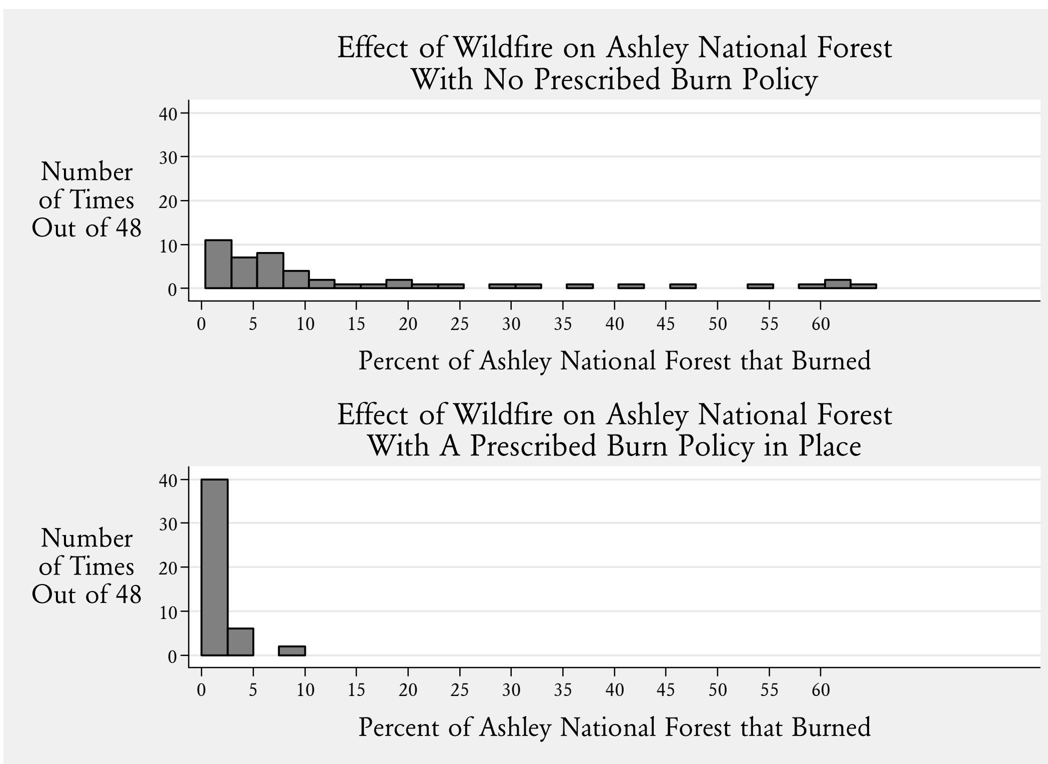

**Figure 2. Illustrative Images from VR Interface**

**Table 2. Betting Task with Stake of $1**

Assume subject has a personal belief that A will occur with probability 0.75

| Bet on A and… | | | | Bet on B and… | | | | Gross expected value of betting | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | A occurs | | B occurs | | A occurs | | B occurs | on A | B | |
| 0.75 | $10 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.11 | 7.50 | 0.28 | 7.22 |
| 0.75 | $5 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.25 | 3.75 | 0.31 | 3.44 |
| 0.75 | $3.33 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.43 | 2.50 | 0.36 | 2.14 |
| 0.75 | $2.5 | 0.25 | $0 | 0.75 | $0 | 0.25 | $1.67 | 1.88 | 0.42 | 1.46 |
| 0.75 | $2 | 0.25 | $0 | 0.75 | $0 | 0.25 | $2 | 1.50 | 0.50 | 1.00 |
| 0.75 | $1.67 | 0.25 | $0 | 0.75 | $0 | 0.25 | $2.5 | 1.25 | 0.63 | 0.63 |
| 0.75 | $1.43 | 0.25 | $0 | 0.75 | $0 | 0.25 | $3.33 | 1.07 | 0.83 | 0.24 |
| 0.75 | $1.25 | 0.25 | $0 | 0.75 | $0 | 0.25 | $5 | 0.94 | 1.25 | -0.31 |
| 0.75 | $1.11 | 0.25 | 40 | 0.75 | $0 | 0.25 | $10 | 0.83 | 2.50 | -1.67 |

**Table 3. Standard Lottery Task in Gains Domain**

| Lottery A | | | | Lottery B | | | | $EV^A$ | $EV^B$ | Difference |
|---|---|---|---|---|---|---|---|---|---|---|
| p($24) | | p($26) | | p($1) | | p($50) | | | | |
| 1 | 24 | 0 | 26 | 1 | 1 | 0 | 50 | 24 | 1 | 23 |
| 0.9 | 24 | 0.1 | 26 | 0.9 | 1 | 0.1 | 50 | 24.2 | 5.9 | 18.3 |
| 0.8 | 24 | 0.2 | 26 | 0.8 | 1 | 0.2 | 50 | 24.4 | 10.8 | 13.6 |
| 0.7 | 24 | 0.3 | 26 | 0.7 | 1 | 0.3 | 50 | 24.6 | 15.7 | 8.9 |
| 0.6 | 24 | 0.4 | 26 | 0.6 | 1 | 0.4 | 50 | 24.8 | 20.6 | 4.2 |
| 0.5 | 24 | 0.5 | 26 | 0.5 | 1 | 0.5 | 50 | 25 | 25.5 | -0.5 |
| 0.4 | 24 | 0.6 | 26 | 0.4 | 1 | 0.6 | 50 | 25.2 | 30.4 | -5.2 |
| 0.3 | 24 | 0.7 | 26 | 0.3 | 1 | 0.7 | 50 | 25.4 | 35.3 | -9.9 |
| 0.2 | 24 | 0.8 | 26 | 0.2 | 1 | 0.8 | 50 | 25.6 | 40.2 | -14.6 |
| 0.1 | 24 | 0.9 | 26 | 0.1 | 1 | 0.9 | 50 | 25.8 | 45.1 | -19.3 |

**Table 4. Standard Lottery Task in the Loss Domain with Initial Credit of $50**

| Lottery A | | | | Lottery B | | | | $EV^A$ | $EV^B$ | Difference |
|---|---|---|---|---|---|---|---|---|---|---|
| p(-$26) | | p(-$24) | | p(-$49) | | p(-$0) | | | | |
| 1 | -26 | 0 | -24 | 1 | -49 | 0 | -0 | -26 | -49 | 23 |
| 0.9 | -26 | 0.1 | -24 | 0.9 | -49 | 0.1 | -0 | -25.8 | -44.1 | 18.3 |
| 0.8 | -26 | 0.2 | -24 | 0.8 | -49 | 0.2 | -0 | -25.6 | -39.2 | 13.6 |
| 0.7 | -26 | 0.3 | -24 | 0.7 | -49 | 0.3 | -0 | -25.4 | -34.3 | 8.9 |
| 0.6 | -26 | 0.4 | -24 | 0.6 | -49 | 0.4 | -0 | -25.2 | -29.4 | 4.2 |
| 0.5 | -26 | 0.5 | -24 | 0.5 | -49 | 0.5 | -0 | -25 | -24.5 | -0.5 |
| 0.4 | -26 | 0.6 | -24 | 0.4 | -49 | 0.6 | -0 | -24.8 | -19.6 | -5.2 |
| 0.3 | -26 | 0.7 | -24 | 0.3 | -49 | 0.7 | -0 | -24.6 | -14.7 | -9.9 |
| 0.2 | -26 | 0.8 | -24 | 0.2 | -49 | 0.8 | -0 | -24.4 | -9.8 | -14.6 |
| 0.1 | -26 | 0.9 | -24 | 0.1 | -49 | 0.9 | -0 | -24.2 | -4.9 | -19.3 |

**Figure 3. Normal and Logistic Cumulative Density Function**

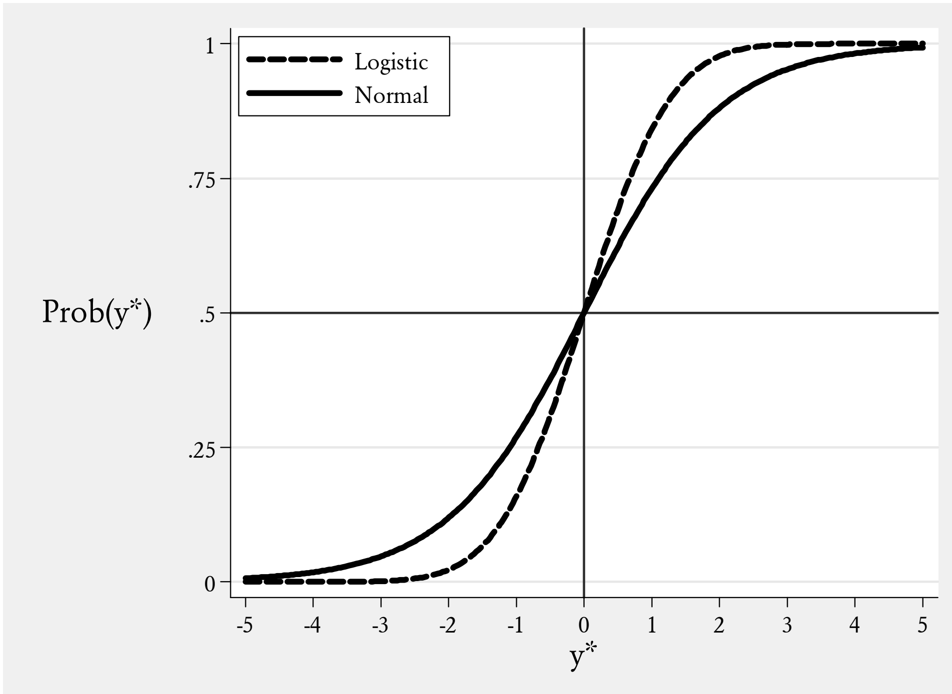


**Table 5. Maximum Likelihood Estimate of Risk Attitudes**

| Parameter | Variable | Point Estimate | Standard Error | *p*-value | Lower 95% Confidence Interval | Upper 95% Confidence Interval |
|---|---|---|---|---|---|---|
| $r$ | Constant | 0.32 | 0.35 | 0.36 | -0.37 | 1.01 |
| | Loss | 0.03 | 0.16 | 0.84 | -0.28 | 0.34 |
| | Betfirst | 0.20 | 0.22 | 0.36 | -0.23 | 0.62 |
| | Age | 0.02 | 0.02 | 0.20 | -0.01 | 0.06 |
| | Female | -0.17 | 0.18 | 0.34 | -0.52 | 0.18 |
| | Hispanic | 0.04 | 0.33 | 0.91 | -0.62 | 0.69 |
| | Business | -0.14 | 0.19 | 0.45 | -0.51 | 0.23 |
| | GPAhigh | 0.33 | 0.27 | 0.23 | -0.21 | 0.87 |
| | Works | -0.56 | 0.26 | 0.03 | -1.07 | -0.06 |
| $\mu$ | Constant | -1.29 | 0.20 | 0.00 | -1.69 | -0.89 |

**Table 6. Joint Estimate of Risk Attitudes and Subjective Beliefs**
**Assuming Homogeneity Across Exogenous and Endogenous Risk Tasks**

| Parameter | Variable | Point Estimate | Standard Error | *p*-value | Lower 95% Confidence Interval | Upper 95% Confidence Interval |
|---|---|---|---|---|---|---|
| $r$ | Constant | 0.20 | 0.25 | 0.43 | -0.30 | 0.70 |
| | Age | 0.01 | 0.02 | 0.53 | -0.03 | 0.05 |
| | Female | -0.14 | 0.15 | 0.34 | -0.44 | 0.15 |
| | Hispanic | -0.01 | 0.20 | 0.94 | -0.41 | 0.39 |
| | Business | -0.26 | 0.17 | 0.12 | -0.59 | 0.07 |
| | GPA high | 0.32 | 0.22 | 0.15 | -0.11 | 0.75 |
| | Works | -0.41 | 0.23 | 0.04 | -0.81 | -0.02 |
| pHI | Constant | -1.16 | 0.34 | 0.00 | -1.82 | -0.50 |
| | Age | 0.05 | 0.02 | 0.02 | 0.00 | 0.09 |
| | Female | 0.43 | 0.47 | 0.37 | -0.50 | 1.35 |
| | Hispanic | 0.55 | 0.37 | 0.13 | -0.16 | 1.27 |
| | Business | -0.87 | 0.43 | 0.04 | -1.71 | -0.02 |
| | GPA high | 0.12 | 0.45 | 0.79 | -0.77 | 1.01 |
| | Works | -0.18 | 0.53 | 0.73 | -1.22 | 0.85 |
| pLO | Constant | 15.48 | 10.46 | 0.14 | -5.02 | 35.99 |
| | Female | -0.71 | 0.94 | 0.45 | -2.55 | 1.14 |
| | Hispanic | 14.35 | 10.85 | 0.19 | -6.92 | 35.61 |
| | Business | -1.11 | 1.57 | 0.48 | -4.20 | 1.97 |
| | GPA high | -13.11 | 10.01 | 0.19 | -32.72 | 6.51 |
| | Works | -0.29 | 1.17 | 0.80 | -2.60 | 2.00 |
| $\mu$ | Constant | -0.94 | 0.19 | 0.00 | -1.31 | -0.57 |
| | Risk | -0.29 | 0.31 | 0.35 | -0.90 | 0.32 |

**Table 7. Joint Estimate of Risk Attitudes and Subjective Beliefs Allowing Heterogeneity Across Exogenous and Endogenous Risk Tasks**

| Parameter | Variable | Point Estimate | Standard Error | *p*-value | Lower 95% Confidence Interval | Upper 95% Confidence Interval |
|---|---|---|---|---|---|---|
| $r$ | Constant | 0.17 | 0.25 | 0.50 | -0.32 | 0.67 |
| | WTP | -0.23 | 0.56 | 0.68 | -1.33 | 0.87 |
| | Age | 0.01 | 0.02 | 0.45 | -0.02 | 0.05 |
| | Female | -0.16 | 0.14 | 0.26 | -0.44 | 0.12 |
| | Hispanic | -0.12 | 0.19 | 0.55 | -0.50 | 0.26 |
| | Business | -0.21 | 0.15 | 0.16 | -0.50 | 0.08 |
| | GPA high | 0.43 | 0.24 | 0.08 | -0.48 | 0.91 |
| | Works | -0.40 | 0.19 | 0.03 | -0.78 | -0.03 |
| pHI | Constant | 1.06 | 0.27 | 0.00 | -1.60 | -0.53 |
| | WTP×age | 0.00 | 0.03 | 0.99 | -0.06 | 0.06 |
| | WTP×female | -0.72 | 0.81 | 0.38 | -2.31 | 0.88 |
| | WTP×hispanic | -0.11 | 0.98 | 0.91 | -2.04 | 1.81 |
| | WTP×business | -2.48 | 1.56 | 0.11 | -5.54 | 0.58 |
| | WTP×GPA high | -0.09 | 0.67 | 0.90 | -1.40 | 1.22 |
| | WTP×works | 0.87 | 0.63 | 0.17 | -0.36 | 2.11 |
| | Age | 0.03 | 0.02 | 0.08 | -0.00 | 0.06 |
| | Female | 0.51 | 0.25 | 0.04 | 0.01 | 1.00 |
| | Hispanic | 0.27 | 0.08 | 0.32 | -0.27 | 0.82 |
| | Business | 0.25 | 0.31 | 0.41 | -0.35 | 0.85 |
| | GPA high | 0.12 | 0.36 | 0.75 | -0.60 | 0.83 |
| | Works | -0.12 | 0.27 | 0.66 | -0.64 | 0.40 |
| pLO | Constant | 15.23 | 0.97 | 0.00 | 13.33 | 17.13 |
| | WTP×age | 0.79 | 0.42 | 0.06 | -.03 | 1.61 |
| | WTP×female | -3.07 | 1.30 | 0.02 | -5.61 | -0.53 |
| | WTP×hispanic | 0.00 | 1.53 | 1.00 | -3.00 | 3.00 |
| | WTP×business | -1.82 | 1.40 | 0.19 | -4.56 | 0.91 |
| | WTP×GPA high | 0.22 | 1.51 | 0.88 | -2.75 | 3.19 |
| | WTP×works | 4.24 | 1.73 | 0.01 | 0.84 | 7.64 |
| | Age | 0.01 | 0.02 | 0.69 | -0.03 | 0.04 |
| | Female | -0.45 | 0.28 | 0.11 | -1.00 | 0.11 |
| | Hispanic | 14.34 | . | . | . | . |
| | Business | -0.96 | 0.61 | 0.12 | -2.16 | 0.25 |
| | GPA high | -13.36 | . | . | . | . |
| | Works | -0.47 | 0.47 | 0.31 | -1.40 | 0.44 |
| $\mu$ | Constant | -1.10 | 0.17 | 0.00 | -1.44 | 0.76 |
| | Risk | -0.11 | 0.30 | 0.72 | -0.70 | 0.48 |

**Figure 4. Estimated Subjective Probabilities of House Burning if Homogeneous Beliefs and Risk Attitudes Assumed**

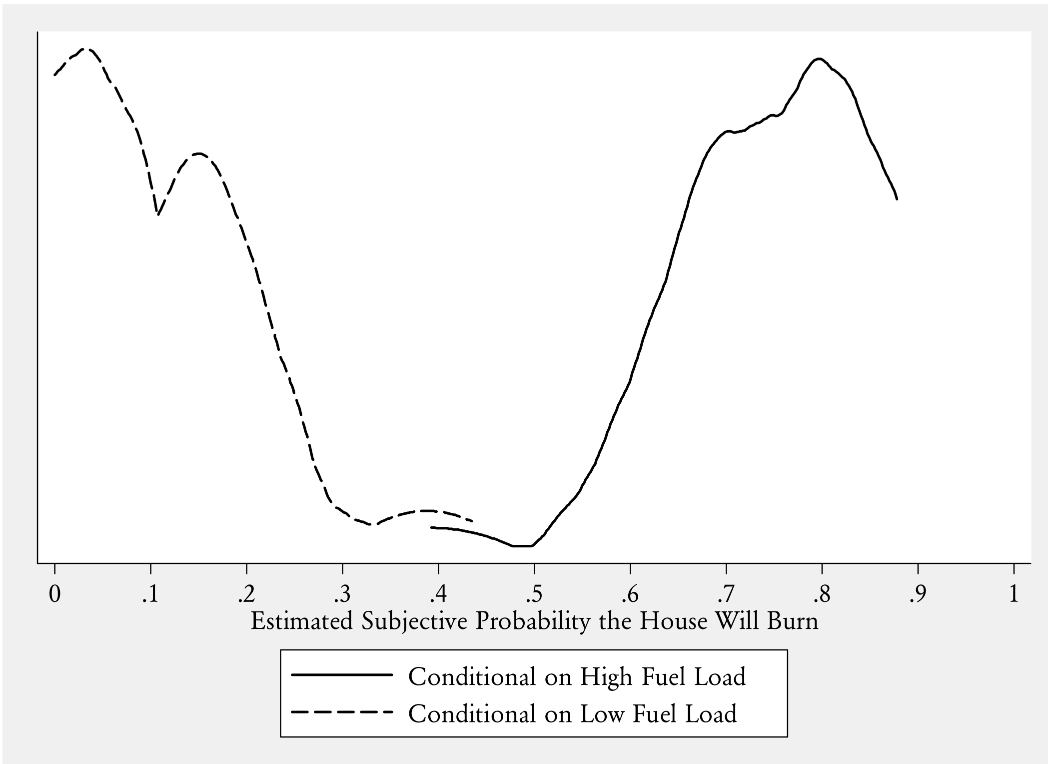

**Figure 5. Estimated Subjective Probabilities of House Burning if Heterogeneous Beliefs and Risk Attitudes Allowed**

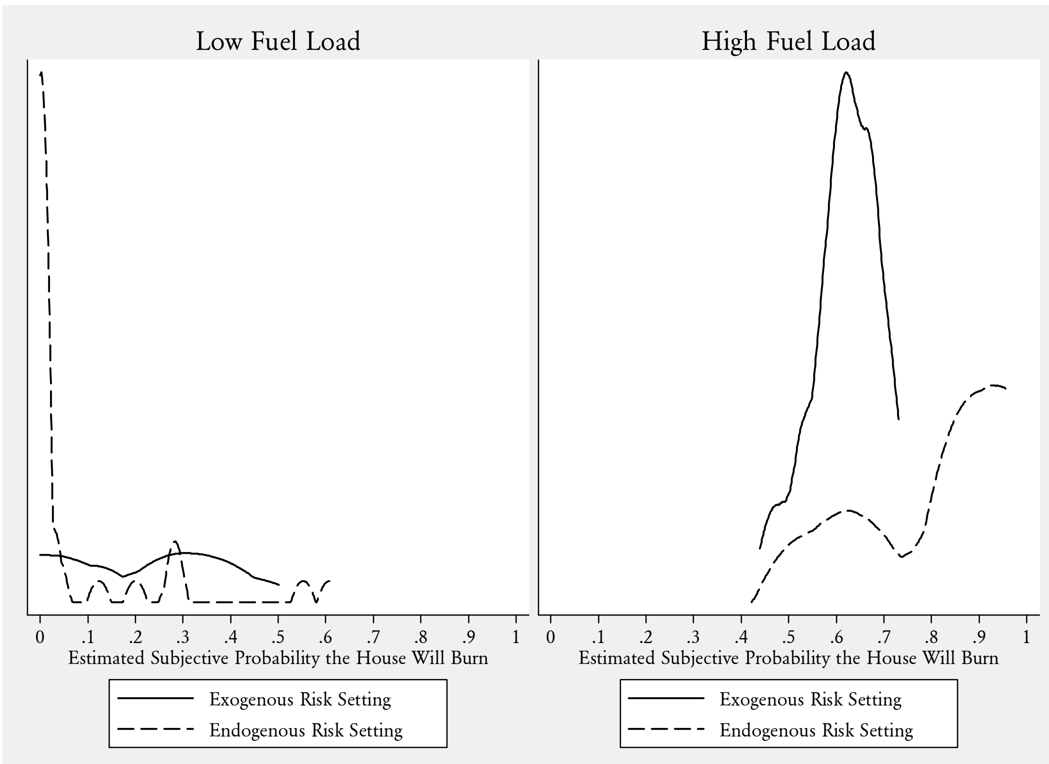

## References


Andersen, Steffen; Fountain, John; Harrison, Glenn W., and Rutström, E. Elisabet, "Estimating Subjective Probabilities," *Working Paper 09-01*, Department of Economics, University of Central Florida, 2009a.

Andersen, Steffen; Fountain, John; Harrison, Glenn W., and Rutström, E. Elisabet, "Estimating Aversion to Uncertainty," *Working Paper 09-04*, Department of Economics, University of Central Florida, 2009b.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten Igel, and Rutström, E. Elisabet, "Eliciting Risk and Time Preferences," *Econometrica*, 76(3), May 2008, 583-618.

Ehrlich, Isaac, and Becker, Gary S, "Market Insurance, Self-Insurance and Self-Protection," *Journal of Political Economy,* 80, 1972, 623-48.

Fiore, Stephen M.; Harrison, Glenn W.; Hughes, Charles E., and Rutström, E. Elisabet "Virtual Experiments and Environmental Policy," *Journal of Environmental Economics and Management*, 57, 2009, 65-86.

Garen, John, "Compensating Wage Differentials and the Endogeneity of Job Riskiness," *The Review of Economics and Statistics,* 70(1), 1988, 9-16.

Harrison, Glenn W.; Humphrey, Steven, and Verschoor, Arjan, "Choice Under Uncertainty: Evidence from Ethiopia, India and Uganda," *Economic Journal*, 119, 2009, forthcoming.

Harrison, Glenn W., and List, John A., "Field Experiments," *Journal of Economic Literature*, 42(4), 2004, 1013–1059.

Harrison, Glenn W.; Johnson, Eric; McInnes, Melayne M., and Rutström, E. Elisabet, "Risk Aversion and Incentive Effects: Comment," *American Economic Review*, 95(3), 2005, 897–901.

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

Harrison, Glenn W., and Rutström, E. Elisabet, "Expected Utility *And* Prospect Theory: One Wedding and A Decent Funeral," *Experimental Economics*, 12(2), June 2009, 133-158.

Harstad, Ronald M., "Dominant Strategy Adoption and Bidders' Experience with Pricing Rules," *Experimental Economics*, 3(3), 2000, 261–280.

Holt, Charles A., and Laury, Susan K., "Risk Aversion and Incentive effects," *American Economic Review*, 92(5), 2002, 1644–1655.

Holt, Charles A., and Laury, Susan K., "Risk Aversion and Incentive Effects: New Data Without Order Effects," *American Economic Review*, 95(3), 2005, 902–912.

Holt, Charles A., and Laury, Susan K., "Further Reflections on Prospect Theory," in J. C. Cox and G. W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).

Kahneman, Daniel, and Tversky, Amos, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 1979, 263-291.

Loubergé, Henri, "Developments in Risk and Insurance Economics: the Past 25 Years," in G. Dionne (ed.), *Handbook of Insurance* (Boston: Kluwer, 2000).

McKee, Michael; Berrens, Robert P; Jones, Michael; Helton, Ryan, and Talberth, John. "Using Experimental Economics to Examine Wildfire Insurance and Averting Decisions in the Wildland-Urban Interface," *Society and Natural Resources,* 17, 2004, 491-507.

Quiggin, John, "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3(4), 1982, 323–343.

Rutström, E. Elisabet, "Home-grown Values and the Design of Incentive Compatible Auctions," *International Journal of Game Theory*, 27(3), 1998, 427–441.

Shogren, Jason F., "The Impact of Self-Protection and Self-Insurance on Individual Response to Risk," *Journal of Risk and Uncertainty,* 3, 1990, 191-204.

Shogren, Jason F., and Crocker, Thomas D., "Risk, Self-Protection, and Ex Ante Economic Value," *Journal of Environmental Economics and Management,* 20, 1991, 1-15.

Shogren, Jason F. and Crocker, Thomas D., "Rational Risk Valuation Given Sequential Reduction Opportunities," *Economic Letters,* 44, 1994, 241-248.

Smith, Vernon L., "Measuring Nonmonetary Utilities in Uncertain Choices: the Ellsberg Urn," *Quarterly Journal of Economics*, 83(2), May 1969, 324-329.

Smith, Vernon L., "Experiments with a Decentralized Mechanism for Public Goods Decision," *American Economic Review,* 1980, 548-590.

Talberth, John; Berrens, Robert P; McKee, Michael and Jones, Michael, "Averting and Insurance decisions in the Wildland-Urban Interface: Implications of Survey and Experimental Data for Wildfire Risk Reduction Policy," *Contemporary Economic Policy,* 24(2), 2006, 203-223.

Wilcox, Nathaniel T., "'Stochastically More Risk Averse:' A Contextual Theory of Stochastic Discrete Choice Under Risk," *Journal of Econometrics*, 142, 2009 forthcoming.