Título:
Probabilistic Sophistication, Sources of Uncertainty, and Cognitive Ability: Experimental Evidence

Autores:
Sergio Almeida (FEA/USP)
Marcos Rangel(Duke University)

Abstract:
This article revisits the idea that individuals are not probabilistically sophisticated, that is, that a well behaved subjective probability measure cannot be constructed from their choice behavior in the presence of uncertainty. The conventional view, based on experimental replications of Ellsberg's (1961) work, is that such simplification is not warranted. In these experiments, however, beliefs are shown to be inconsistent across sources of uncertainty that involve known and unknown probabilities. We design and implement an experiment with uniform sources of uncertainty, and find that the majority of subjects satisfy requirements of probabilistic sophistication. We also show that the few violations are strongly correlated with low cognitive ability.

Resumo:
Este artigo revisita a idéia de que os indivíduos não são probabilisticamente sofisticado, isto é, que uma medida bem comportada de probabilidade subjetiva não pode ser construída a partir de decises sob incerteza. A viso convencional, baseada em replicaes do experimento de Ellsberg (1961), a de que essa simplificao no factvel. Todavia, tem sido demonstrado que as crenas probabilsticas nesses experimentos de Ellsberg so inconsistentes entre fontes distintas de incerteza. Nesse artigo, desenhamos um experimento em que as fontes de incerteza so uniformes. Os resultados mostram que nesse contexto, a maioria dos indivduos satisfaz os requisitos de sofisticao probabilstica. Mostramos que as poucas violaes encontradas so fortemente correlacionada com habilidade cognitiva relativamente baixa.

Keywords:
probabilistic sophistication; cognitive ability; decision making under uncertainty

Palavras-chave:
sofisticação probabilística; habilidade cognitiva; decisão sob risco e incerteza

Área ANPEC:
8

JEL Code:
D01; D81; C91

Probabilistic Sophistication, Sources of Uncertainty, and Cognitive Ability: Experimental Evidence

Sergio Almeida, Marcos Rangel

July 2014

Abstract

This article revisits the idea that individuals are not probabilistically sophisticated, that is, that a well behaved subjective probability measure cannot be constructed from their choice behavior in the presence of uncertainty. The conventional view, based on experimental replications of Ellsberg's (1961) work, is that such simplification is not warranted. In these experiments, however, beliefs are shown to be inconsistent across sources of uncertainty that involve known and unknown probabilities. We design and implement an experiment with uniform sources of uncertainty, and find that the majority of subjects satisfy requirements of probabilistic sophistication. We also show that the few violations are strongly correlated with low cognitive ability.

1 Introduction

The relationship between attitudes toward risk and attitudes toward uncertainty has been a central issue in the literature on individual decision making. Interest in the topic stems from earlier theoretical work by Savage (1954). According to Savage's work, choices with consequences that depend on the states of the world for which no exogenous probability exists can be modeled as maximizations of expected utility, with the decision maker's beliefs replacing objective probabilities. Further theoretical formalization provided by Machina & Schmeidler (1992) and Grant (1995) indicate that when beliefs resemble a classical probability measure, preferences over *acts* (a function that maps states to outcomes) mirror preferences over objectively risky *lotteries* constructed using her beliefs. For these *probabilistic sophisticated* agents, choices under uncertainty can be characterized as expected-utility-maximizing behavior just like in standard models of decision making under risk.

The empirical literature on the topic, however, has questioned the validity of the concept of probabilistically sophisticated behavior among economic agents. Ellsberg (1961) and several subsequent replications showed that individuals can display: i) betting preferences that implicitly assign the same probability to two exhaustive events; and ii) strict preference for a bet with a known 50-50 probability distribution for the events over the same bet with unknown probabilities.[1] Nevertheless, to the extent that probabilistic sophistication is a consistency requirement for preferences over symmetric choice objects in two different "worlds" (risk and uncertainty), one might argue that Ellsberg-like experiments contrasting preferences over known and unknown probabilities may be a more stringent test of probabilistic sophistication than necessary. Our review of the literature reveals, however, that

[1] See Kahn & Sarin (1988), Heath & Tversky (1991), Fox & Tversky (1995), Cohen *et al.* (1985), and Halevy (2007). For a more comprehensive list of the first wave of replications and variations on Ellsberg's experiment, see Camerer & Weber (1992).

very little is known empirically about whether probabilistic sophistication is satisfied when events are generated by a common source of uncertainty.[2]

In the present article, we contribute to this debate by demonstrating experimentally that when there is common uncertainty about events, most individuals' preferences do not violate probability sophistication. When adopting the least restrictive violation criteria, for example, we find that less than 13 percent of our subjects have preferences that are inconsistent with probabilistic sophistication. Moreover, the violations that occur are strongly correlated with relatively low cognitive ability.

The conceptual framework that guides our analysis is drawn from work by Grant (1995), Strzalecki (2011), and Karni (2009) and delivers a tractable definition of probabilistic sophistication that imposes no restrictions on specific functional representations.[3] In our case, probabilistic sophistication predicts that preferences over lotteries that are built up with the individual's subjective probabilities and defined over the set of outcomes of the acts will exhibit the same preference ordering over the acts within a given source of uncertainty (one in which vagueness remains information-wise symmetric).[4] Building on this prediction, we design an experiment with a sequentially independent three-part procedure in which we elicit each subject's preferences over subjectively uncertain acts, their risk aversion, and the subjective probability distribution over all events in the state space of the initial choice problem.[5] By measuring subjects' attitudes toward risk and their probabilistic beliefs in an incentive-compatible fashion, we can perform a simple test to assess the extent to which choice behavior is probabilistically sophisticated. We then extend the analysis in order to explore whether behavior that violates probabilistic sophistication is correlated with cognitive ability. In our experiment, we use tests that measure different cognitive skills to evaluate the robustness of the results to the manifestation of different types of intelligence in the sample. In particular, we implement tests that aim to measure fluid and crystallized aspects of human intelligence. The former is related to subjects' ability to solve new problems, whereas the latter is the accumulated knowledge that is necessary to perform calculations and/or to suppress impulsive responses.[6] According to our results, *ceteris paribus*, the greater an individual's cognitive ability, the more likely the individual is to exhibit probabilistic sophistication. If probabilistic sophistication is understood as a form of rationality, this result suggests that adherence to the normative principles of models of rational decision depends not only on the "environmental" conditions under which decision making occurs (e.g., incentives, learning opportunities) but also on the individual's internal make-up. We provide a simple illustrative model that explores psychologists' dual-process accounts of reasoning, both in order to provide a

[2]Abdellaoui *et al.* (2011) find that violations of probabilistic sophistication occurred for comparisons of references across different sources of uncertainty but not within them. The researchers define probability sophistication as preferences over bets (mappings from events to outcomes) that depend exclusively on the probability distribution over the outcomes (p. 698) and are not affected by the positioning of the events within the partitions within sources. In addition to not being the definition widely used in this literature (see Machina & Schmeidler (1992), Grant (1995), Marinacci (2002), Karni (2009) and Strzalecki (2011)), this makes their finding somewhat unsurprising, since it boils down to a test of event partitions' exchangeability within a given source.

[3]This framework is consistent with the approach adopted by Ramsey (1926), Savage (1954),Anscombe & Aumann (1963), and Machina & Schmeidler (1992) demonstrating that the characterization of probabilistic beliefs can be separated from restrictions on preferences.

[4]This is closely related but not necessarily identical to what Abdellaoui *et al.* (2011) call a uniform source of uncertainty. Those authors' definition of uniformity requires events to be exchangeable within the source in addition to requiring the symmetry of the probability information. For our purposes, exchangeability is not a requirement.

[5]By design, we only deal with additive probability measures. Models in which the decision maker bases his choices on beliefs that can be represented by a single additive probability have been used since Ramsey (1931) and de Finetti (1937). These models provide a framework that a wide variety of economic models, including models of expected and non-expected utility, use to deal with real-world uncertainty that does not present itself in terms of exogenously given probabilities. See, for instance, Pratt *et al.* (1964), Anscombe & Aumann (1963), Fishburn (1969), Schmeidler (1989), Gilboa & Schmeidler (1989).

[6]Regarding fluid and crystallized forms of intelligence, see Horn and Horn & Cattell (1967). For "cognitive reflection", see Frederick (2005).

theoretical support to our second result and also to suggest directions for further scientific inquiry on how "rational" choice behavior can be bounded by cognitive ability.

The remainder of the paper is organized as follows. Section 2 establishes theoretical preliminaries that are essential to impose conditions on preferences that exhibit probabilistic sophistication. Section 3 describes the experimental design and procedures. Section 4 presents the theoretical predictions for choice behavior in our experiment. The test methodology is described in Section 5, while Section 6 reports and discusses the results. Section 7 sketches a simple model of dual reasoning that can account for the pattern of violations observed across individuals in our sample. Conclusions are presented in Section 8.

2 Characterizing Probabilistic Sophistication

Probabilistic sophistication is informally defined as choice decision based on probabilistic beliefs (Strzalecki, 2011, Abdellaoui *et al.*, 2011). However, probabilistic sophistication actually imposes more structure on preferences than this informal definition might suggest; in fact, it is a consistency condition for preferences over state-contingent bets (uncertainty) and preferences over lotteries (risk) based on the distribution of the individual's probabilistic beliefs about states over the same set of outcomes. To examine this condition, consider a setting in which individuals have preference relations over pairs of acts or lotteries.

ACTS: Acts are bets that map the states of nature to monetary outcomes. Let $S = \{s, t, ...\}$ be a finite set of states, one of which is the true state. Subsets of S are events; let $\mathcal{E}$ be a σ-algebra of all subsets of S. Let $M = \{m : m \in [a, b]\, b > a\}$ be a finite set of monetary outcomes. $A = \{f, g, h, ...\}$ is a set of simple two-consequence acts: finite-valued mappings from S to M that attach a monetary outcome to each possible state. An act f_m is termed constant if it assigns the same monetary outcome m to every state. Let $\succsim$ be a preference relation defined on the set of acts A. $\succ$ and $\sim$ correspond to the strict and indifference preference relations, respectively.

LOTTERIES: Lotteries are probability distributions over a set of outcomes. Let $L(M) = \{\ell(x, p, y, 1-p) : x, y \in M, p \geq 0, y > x\}$ be the set of all binary lotteries over the set $x, y \in M$ $(x > y)$ of outcomes. For the sake of simplicity, we shall write $\ell(x, p, y)$ to denote the lottery that pays x with probability p and y with probability $1 - p$. Let $l(x, 1)$ stand for the degenerate lottery in which outcome x is certain and the other outcome has a probability of zero. Let $\succsim_L$ be a preference relation defined for the set of binary lotteries $L(M)$. $\succ_L$ and $\sim_L$ correspond to the strict and indifference preference relations, respectively.

Suppose now that the decision maker has a finitely additive nonatomic probability measure μ on $\mathcal{E}$ implicit in $\succsim$. This probability measure can then be used to map each act in A to a lottery in $L(M)$. For example, a simple act h that pays x if the event $E \in \mathcal{E}$ obtains and y otherwise would be mapped to the lottery $\ell(\mu(E), x, y)$ as $h^{-1}(x) = E$, $h^{-1}(y) = S/E$. This probability measure then allows preferences over acts to induce preferences over lotteries: if μ maps simple acts g and h to binary lotteries ℓ and ℓ^*, respectively, then $g \succsim h$ implies $\ell \succsim_L \ell^*$. Note that this subjective probability measure allows the decision maker to treat subjective uncertainty as risk and rank acts and lotteries on the basis of their implied probability distributions over outcomes. This is precisely the principle underlying the notion of probabilistic sophistication.

DEFINITION 1: *A decision maker with a preference relation* $\succsim$ *on* A *exhibits probabilistic sophistication if his probability measure* μ *on* E *induces a preference relation over the lotteries in* $L(M)$ *such that any simple act* f *in* A *and lottery* $\ell(x, p, y)$ *in* $L(M)$, $p \in [0, 1]$, $\mu(f^{-1}(x)) = p$ *implies* $f \sim \ell(x, p, y)$.

Using this definition, we can assess the outcomes of an act according to the likelihood of the events through which each outcome arises, not the nature or source of the events themselves. As shown in a more general setting by Grant (1995), Strzalecki (2011) and others, one can reconstruct the preferences of a probabilistically sophisticated decision maker based on his beliefs about the likelihood of events associated to each of the outcomes in question and based on the decision maker's preferences over probability distributions given equivalent outcomes. Building on this idea, we design an experiment that tests for probabilistic sophistication and examine whether any violations can be explained by bounded cognition.

3 Experimental Design and Procedures

Our experimental design consists of distinct individual tasks that address the goals of (1) measuring the subjects' cognitive ability, (2) eliciting their preferences over uncertain bets, (3) determining their preference over lotteries involving the same distribution of bet outcomes, (4) reconstructing their beliefs about the likelihood of events that will determine the payoff of the uncertain bets, and (5) eliciting an alternative measure of the subjects' cognitive ability and identifying some of their socio-demographic characteristics. We implemented this design in a five-part procedure.

Part one was designed to measure the subjects' cognitive ability. We asked the subjects to complete a "fluid intelligence test" based on Raven's *Progressive Matrices*. This test measures one's ability to think abstractly, to identify patterns in novel problems and to generate meaning from confusion. Although the test is by no means a measure of general cognitive ability, which is influenced by very different types of knowledge and motivation levels, fluid intelligence includes numerous problem-solving skills and is one of the best single measures of g as reported by several studies (for a survey, see Carrol, 1993). The test consists of a set of twelve questions, each of which features a figure with symbolic patterns presented in the form of a 3×3 matrix with one symbol missing. Viewing the set of symbols, the subject must identify the missing symbol that is required to complete the pattern in the figure. Test writers chose symbols on test questions that attempt to minimize the score contamination that might occur due to interpersonal variations in familiarity with symbols that have cultural meaning (Raven, 1998). For the sake of variation, we limited the amount of time that the subjects had to complete the test to only three minutes. This limitation made our test different from Raven's Progressive Matrices. An individual's test score is the number of questions answered correctly.

Part two presented the subjects with a discrete choice involving uncertainty. The subjects were asked to choose between two options involving two states of nature, s and t. Option A was a constant act that resulted in a definite outcome of 20 tokens in every state, whereas option B yielded a higher outcome of 45 tokens in state s and a lower outcome of 10 tokens in state t. Option B was an uncertain two-consequence act whose uncertainty is resolved in two stages after choices are made. In the first stage, an observation is obtained, and the probability of each state (where $p_s + p_t = 1$) is defined. In the second stage, the state is realized, and the payoff of option B accrues. The subjects were informed when they made their choices that the probability distribution over states would be defined according to a simple function g that mapped each event in E to a state probability distribution as follows: $g(e_i) = [1 - e_i/N] = p_s$. Here, the expected value of act B is increasing in the outcome of the event that obtains.[7] We contend that the advantage of this design (in which a choice is made

[7] As part of a larger study, we implemented two treatments that varied the source of uncertainty over the state probabilities by manipulating the nature of the events. In one treatment, the events and therefore the state probabilities were defined by the distribution of actions chosen by the players such that p_s (p_t) is the proportion of N subjects who chose the safe (uncertain) act. In the other treatment, the events were defined by the distribution of orange and green balls in an urn containing N balls of either color; subjects were told that a machine that followed an unknown process selected the distribution. In this case, p_s (p_t) was the proportion of orange (green) balls. From a descriptive point of view, there is no reason to believe that the subjective mapping of events into state probability distributions will systematically differ between the two treatments, as predicted by most models of decision making under uncertainty. Nonetheless, our data analysis controls for potential treatment effects that might affect the degree to which a subject's choice in this step is consistent with her subjective probability distribution over states of nature (elicited in part four).

between a degenerate act and an uncertain one) over a more general design (in which a choice is made between a pair of purely non-degenerate acts and lotteries) is that the chosen design simplifies the choice problems and the subjective assessment of outcome-relevant events without compromising the representation of probabilistically sophisticated preferences in this context.

Part three elicited the subjects' levels of uncertainty aversion. We followed Bohnet *et al.* (2008) by asking the subjects to state the smallest probability p of receiving 45 tokens for which they would choose the uncertain act over the degenerate one with a certain outcome of 20 tokens. A subject's stated "minimum acceptable probability" (henceforth, MAP) was binding in the sense that it was used to select the actual action in the above-mentioned step that, in turn, determined her payoff at the end after uncertainty over p_s is resolved. If her MAP was larger than the actual p_s in her session (let that be $\hat{p}$), the subject was then assumed to opt for option A and received the sure outcome of 20. If her MAP was less than or equal to $\hat{p}$, she was assumed to select the uncertain act, which left her with a $\hat{p}$-chance of receiving 45 tokens and a $(1-\hat{p})$-chance of receiving 10 tokens. As long as an individual believes that $\hat{p}$ can take values within an arbitrarily close neighborhood around her MAP and conform to the substitution axiom from expected utility theory, her reported MAP will be incentive compatible. Therefore, this procedure elicits the true "least favorable" probability distribution over option B's two outcomes, 45 and 10, that the individual regards as better than the sure outcome of option A, 20. Hence, from a subject's MAP, we can infer all of the probability distributions over 45 and 10 tokens that are strictly preferred to the option of 20 tokens. Later, we shall see that these probability distributions, in conjunction with an individual's expected subjective probability distribution over the states of nature, fully characterize the preferences regarding the choice options in part two that a probabilistically sophisticated decision maker should exhibit.

In the *fourth part* of the experiment, the subjects were asked to report their beliefs about the likelihood of each event in E, an integer number between 0 and 10 that straightforwardly defines the probability distribution over the two prizes of option B. The subjects were not aware that they would be participating in this belief elicitation procedure when they completed the previous portions of the study. They were given instructions about how to state their probabilistic beliefs and were told how they would be rewarded for accurate statements. More specifically, the subjects were asked to state their probabilistic beliefs for each possible "event" from part two. The subject's stated belief about the likelihood of each possible event could be represented using any integer number between 0 and 100. We imposed additivity by design, giving the subjects 100 poker chips to be allocated among the events according to the subjects' subjective beliefs about the likelihood of each event. Each event transparently maps onto the a posteriori probability measure for the "good" outcome of option B. Thus, given the additivity restriction, the subjects' likelihood estimates can be reasonably interpreted as their complete subjective probability distribution over $\hat{p}$, the actual probability of earning 45. To encourage the subjects to provide truthful and accurate reports, we used a proper scoring rule to reward accurate belief statements. Although the incentive compatibility of this rule relies on restrictive assumptions (see, e.g., Offerman *et al.*,2009), the proper scoring rule has been shown to produce estimates that are no different from those obtained using other methods - including risk aversion-corrected results. These estimates also are not subject to the hedging-like behavior of risk-averse individuals (Trautmann & van de Kuilen, 2010, Blanco *et al.*, 2010).

The proper scoring rule involves a quadratic loss function defined as follows. Let subject i's stated beliefs regarding the likelihood of each of the N events that define the state probabilities be p_i^{Ω}, a probability distribution over the set $(e_1, e_2, ..., e_N)$, that is $p_i^{\Omega} \equiv \{p_i^{e_1}, p_i^{e_2}, ..., p_i^{e_N}\}$, such that $\bigcup_{i=1}^{N} e_i \equiv \Omega$ and $p_i^{\Omega} \in \Theta \equiv \{p_i^{e_k} \in [0,1] | \sum_{i=1}^{N} p_i^{e_k} = 1\}$. Once the uncertainty is resolved, let $\hat{p}_{\Omega} \equiv \{\hat{p}_{e_1}, \hat{p}_{e_2}, ..., \hat{p}_{e_N}\}$ be the indicator of whether an event occurs, such that $\hat{p}_{e_k}$ equals 1 if the event e_k occurs and equals 0 otherwise. The quadratic scoring rule then determines that subject i's payoff from his probabilistic belief distribution will be $\pi_i(p_i^{\Omega}, \hat{p}_{\Omega}) = a - b\sum_{k=1}^{N}(p_i^{e_1} - \hat{p}_{e_k})^2$ where a and b are constants. We set $a = 5,000$ and $b = 0.25$ so that using $1,000$ points as a unit of payment would produce integer scores with four digits of precision. Thus, the maximum score for one question was

5,000, the minimum score was 0, and the certain score that resulted from the distribution of poker chips that was as close as possible to a uniform distribution was 2,772.50 points. In this part of the experiment, the sum of the points for each subject was converted to money using an exchange rate of 1,000 points = R$ 1.

In the *fifth part* and last part of the experiment, the subjects were asked to complete a short demographic questionnaire. Frederick's (2005) cognitive reflection test (CRT) was included in this questionnaire to complement our cognitive ability measure. The CRT is a simple three-item test that is designed to measure one's ability to control impulsive thinking: that is, the ability to suppress an erroneous answer to a problem that springs "impulsively" to mind (Frederick, 2005, p.27). Because supplying correct answers on this test requires learning and previously acquired knowledge, we believe that the test can be viewed as a measure of *crystallized intelligence*. By using the subjects' scores on this test and the *fluid intelligence* test described above, we provide the measures of both components of the Cattel-Horn model of intelligence (Cattell, 1943, Carrol, 1993).

3.1 Administration

A pool of 197 subjects was recruited from the graduate and undergraduate student body at the University of Sao Paulo in Brazil. The subjects were randomly assigned to 20 experimental sessions with N=10 subjects each.[8]

The experiment was conducted using a paper-and-pencil format, with sessions that lasted no more than 70 minutes. Upon their arrival, the subjects sat at visually isolated desks. The instructions for each part of the experiment were distributed separately and were read aloud by the experimenter as the subjects read along on paper. After all parts of the experiment had been completed, we resolved the uncertainty in two stages. First, we resolved the uncertainty about which event occurred, obtaining an integer number N between 0 and 10. The outcome was used to calculate $\hat{p}$ $(= N/10)$ and $(1 - \hat{p})$, the actual probability distribution over the two consequences of act B. We then compared $\hat{p}$ with each subject's MAP to decide which act in part two would be used to determine the subject's earnings for that part of the experiment. If $\hat{p} < MAP$, the subject selected the constant act and earned 20 tokens. Otherwise, the subject selected the uncertain act, which we resolved by drawing a ping-pong ball from an urn containing $\hat{p} \times 10$ oranges and $(1 - \hat{p}) \times 10$ greens. If the ball drawn was orange, the subjects earned the "good" prize of 45 tokens; otherwise, they earned the "bad" prize of 10 tokens. At the end, each token earned in the second part of the experiment was converted into money using an exchange rate of 1 token = 1 Brazilian Real (at the time, R$1,00 = US$ 0,64). Those earnings were added to the earnings from the belief elicitation process conducted during the fourth part of the experiment, along with a R$5 show-up fee. The total average payment was R$ 29.09 per participant.

4 Theoretical Prediction

Here, we outline a decision setting that mimics the one in our experiment. For a given subjective probability distribution and preference over lotteries, we derive the consequences of probabilistic sophistication for preferences over acts.

Let g be an act that pays x if the true state is s and y if the true state is not s. Let h_z be a degenerate act that pays z in every state, where $x > z > y$. Let $I = \big\{i : i \in \{0, 1, 2, ..., N\}\big\}$ be a finite set of informational events, and let $L = \{\ell(x, p, y) : 0 \leq p \leq 1, x, y \in \mathbb{R}_+\}$ be the set of all binary lotteries that pay x with probability p and y with probability $1 - p$. Let $P : I \times (x, y) \to L$ be a function that maps each informational event and each pair $(x, y) \in \mathbb{R}^2_+$ of outcomes onto a

[8]We were able to recover undergraduate transcripts using the university's electronic system. We will use these data to control for differences in statistical training.

particular lottery $l \in L$ in the following way: $P(i,x,y) = \ell(x, 1-i/N, y)$. Each informational event assigns an objective probability to state s and therefore to the outcome x. Thus, the uncertainty surrounding the informational event represents the uncertainty surrounding the true state of nature.

Suppose, then, that the decision maker has an additive probability measure $\mu(.)$ that is defined on I so that $\mu(i)$ denotes the probability that the decision maker assigns to the informational event i. If we know the decision maker's subjective probability distribution over the informational events, we can use this information to map acts g and h_z to binary lotteries in the following way: $g = \left\{\ell(x, \mu(s), y) \in L : \mu(s) = \sum_{i \in I} \mu(i) P(i,x,y)\right\}$. h_Z maps straightforwardly to $l(0,x,z)$, a degenerate lottery that pays z with certainty.

Suppose we know the decision maker's MAP. Recall that MAP is a probability $p \in [0,1]$ such that $\ell(x,p,y) \succsim \ell(z,1)$ for all $p \geq MAP$. From one's MAP, we can infer $L_{\succ_z} = \{\ell(x,p,y) \in L : p \geq MAP, x > z > y\}$, the subset of the binary lotteries in L that are preferred over the degenerate lottery that pays z with certainty. A graphical illustration of this result is displayed in Figure 4, where the thicker segment of the line linking $u(\underline{x})$ and $u(\bar{x})$, the utility of the "low" and "high" outcomes of the lottery, respectively, represents $L_{\succ_z}$.

Figure 1: MAP: A Expected Utility Interpretation

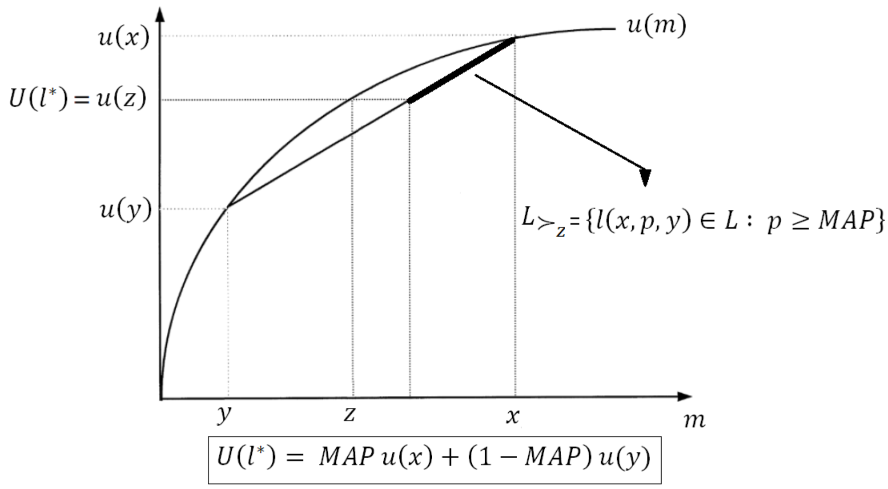

If the decision maker's preferences over lotteries captured by his MAP satisfy the condition of monotonicity, then these preferences in conjunction with $\mu(.)$ fully characterize the preferences over acts g and h_z that a probabilistically sophisticated decision maker ought to have. We summarize this finding in the following prediction, the proof of which appears in the appendix:

PREDICTION: *Probabilistically sophisticated individuals will exhibit one of the following two patterns of choice behavior:*

(P1) $g \succ h_z$ *if* $\mu(s) > MAP$
(P2) $h_z \succ g$ *if* $\mu(s) < MAP$

where $\mu(s)$ *denotes the probability that the decision maker assigns to* s*, the state in which* g *pays the "high" outcome* x*. There is indifference between acts in the case of equality.*

5 Violations of Probabilistic Sophistication: Test Criteria

The above prediction yields a straightforward method of testing for violations of probabilistic sophistication. Depending on whether their MAP is above or below the average expected probability of act g paying x, for probabilistically sophisticated individuals, either $g \succ h_z$ or $h_z \succ g$. Although this threshold-like rule is directly derived from the theoretical definition, the rule might not capture violations of probabilistic sophistication if the decision maker uses probabilistic rules to make his choice. To glean insight into the sensitivity of the observed degree of violation of such rules, we take the liberty of moving beyond the "deterministic" choice rule implicit in the above prediction. Building on this rule, we propose additional probability-based test criteria for characterizing choice patterns in our experiment that violate probabilistic sophistication. This finding will be clarified in the next subsections.

Assuming that an individual's preference over acts is probabilistically sophisticated if, loosely speaking, it is consistent with his beliefs, our test criteria will allow us to explore the implications of the differing levels of strength of the consistency conditions for the observed degree of violation and the relationship between such degree of violation and cognitive ability. In addition, the criteria will help us to test the robustness of our results in the presence of alternative, less "deterministic" and stronger criteria for the violation of probabilistic sophistication.

5.1 "Weak" Violation

First, we consider the baseline case of probabilistic sophistication as stated in the above prediction. In this case, preferences over subjectively uncertain lotteries are constructed from the subjective probability distribution via averaging. These preferences are then contrasted with their objectively uncertain counterparts (inferred from the subject's MAP) to assess whether they are consistent with the subjective probability distribution. We shall refer to this measure as the "reductionist" test criterion for probabilistic sophistication because an individual's probabilistic beliefs about the probability of the higher prize, $p_{\bar{x}}$, are summarized by the first moment of the distribution of such beliefs, $E[p_{\bar{x}}]$.

CRITERION 1 ("Reductionist"): *A preference of g over h_z (h_z over g) violates probabilistic sophistication if $E[p_{\bar{x}}] < MAP$ $\big(E[p_{\bar{x}}] > MAP\big)$.*

Notice that many distributions of beliefs can make a particular preference between g and h_z for a given MAP inconsistent with probabilistic sophistication; therefore, we regard this violation as "weak". As shown in section 2, Criterion 1 remains a natural implication of the very definition of probabilistic sophistication. We can illustrate the "weakness" of the criterion using an example.

Example 1 below illustrates an instance in which, for a given MAP, three different probability distributions of beliefs make the preference for the "safe option", act h_z, over the "uncertain" one, g, inconsistent with probabilistic sophistication.

EXAMPLE 1: Consider the case of three individuals who prefer h_z over g. Let $MAP_i = p^*$ $\forall i = \{1,2,3\}$, and f_i the probability density function of i's beliefs about $p_{\bar{x}}$, the probability of act g yielding x. Let $\mu_n^i = \int_0^1 p_{\bar{x}}^n f(x_{p_{\bar{x}}}) dp_{\bar{x}}$ be the n-moment about zero of the probability distribution of i's beliefs about the probability of obtaining the higher prize associated with option g. Assume that for $n > 1$, $n > 1$, $\mu_n^i \neq \mu_n^j \quad \forall i \neq j$, with equality holding among the distributions of beliefs for $n = 1$, which means that these probability distributions have the same mean but that the mass is distributed very differently among them. Assume $p^* < \mu_1^i \quad \forall i$. The three distributions of beliefs about $p_{\bar{x}}$ then make $h_z \succ g$ inconsistent with the concept of probabilistic sophistication. Figure 2 below graphically illustrates this example.

Figure 2: Subjective PDFs of Higher Prize Probability (For all subjects, $h_z \succ g$).

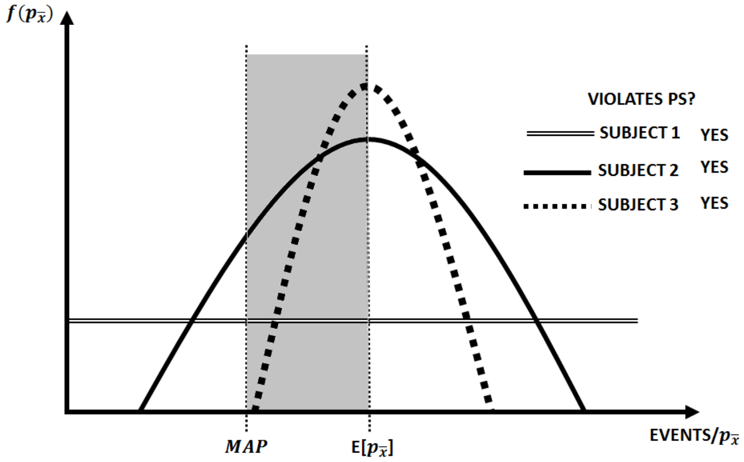

5.2 "Strong" Violations

We now derive criteria for the violation of probabilistic sophistication using the hypothesis that for a given MAP an individual's choice between acts g and h_z will depend on the distribution of his beliefs about the distribution of $p_{\bar{x}}$, the probability of g yielding x. These criteria will capture the possibility that individuals, given the stochastic nature of their probabilistic beliefs, use probabilistic rules to decide which act to choose.

First, we consider the "hypothesis test"-like case in which an individual for whom $MAP = p^*$ will decide whether to choose between g and h_z based on the subjective probability that $p_{\bar{x}} \geq p^*$ is above a certain level θ.

CRITERION 2 ("Hypothesis test-like"): *Let $F(.)$ be the probability distribution function of $p_{\bar{x}}$. A preference of g over h_z (h_z over g) is probabilistically sophisticated if $F\Big(p_{\bar{x}} \geq MAP\Big) > \theta$ $\Big(F\Big(p_{\bar{x}} \geq MAP\Big) \leq \theta\Big)$.*

In the results section, we consider the case in which $\theta = 0.1$. Figure 3 below revisits Example 1 and illustrates how under Criterion 2, not all three probability distributions of beliefs about $p_{\bar{x}}$ will make $h_z \succ g$ inconsistent with probabilistic sophistication for a given MAP that is below the mean of the distribution.

Figure 3: Subjective PDFs of Higher Prize Probability (For all subjects, $h_z \succ g$).

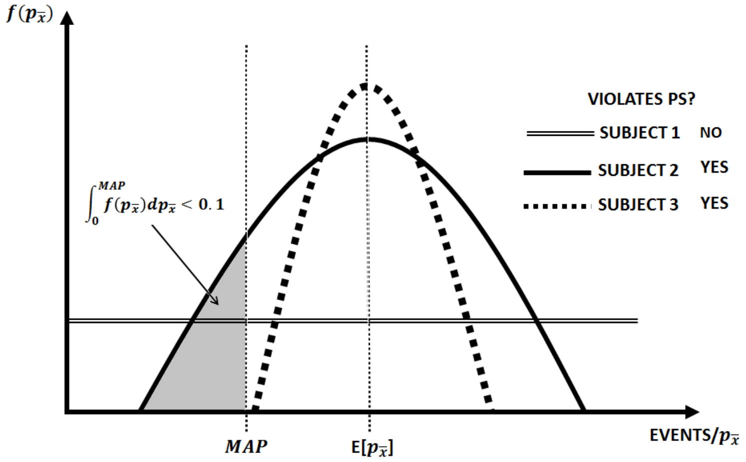

We now consider what we refer to as the "disbelief in beliefs" case, in which an individual with a $MAP = p^*$ will decide between g and h_z based on the subjective probability of $p_{\bar{x}} \geq p^*$ being above zero. This is an extreme case of Criterion 2 in which $\theta \in \{0, 1\}$. The patterns of choices that do not satisfy this criterion are the strongest possible form of the violation of probabilistic sophistication in the sense that there is no probabilistic belief about $p_{\bar{x}}$ to justify the choices made; hence, it is as though the choices were made in complete disregard of the individual's beliefs.

CRITERION 3 ("Disbelief in beliefs"): *Let $F(.)$ be the probability distribution function of $p_{\bar{x}}$. A preference for g over h_z (h_z over g) is probabilistically sophisticated if* $F\Big(p_{\bar{x}} \geq MAP\Big) = 1$ $\Big(F\Big(p_{\bar{x}} \geq MAP\Big) = 0\Big)$.

Figure 4 below revisits Example 1 and illustrates how under Criterion 3, only the choice of the individual with the low kurtosis distribution violates the requirement of probabilistic sophistication because $h_z \succ g$ even though the entire probability mass of the individual's distribution of $p_{\bar{x}}$ is located above MAP.

Figure 4: Subjective PDFs of Higher Prize Probability (For all subjects, $h_z \succ g$).

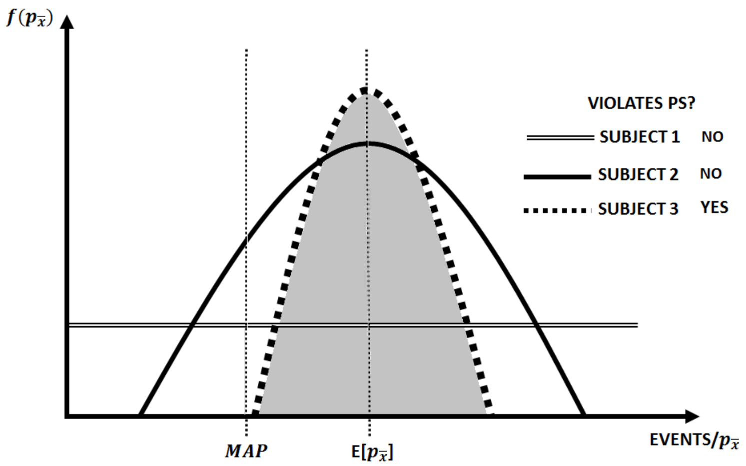

6 Results

Before presenting the results of our experiment, we present descriptive statistics for the pool of participating subjects. Approximately 60 percent of the individuals in the experiment are males and that approximately 70 percent of them are white. The subjects' average age is 23 years old, and most live in households whose heads are college educated. More than half of the subjects attend the university's School of Business, Economics and Accounting (which we classify as part of the social sciences division of the university) and 66 percent of the subjects had already taken at least one course in statistics and/or econometrics.

The average individual has a minimum acceptance probability of 43.4 percent and an expected occurrence of the best outcome of approximately 46 percent. This result indicates that if the average individual were one of the participants, the sophisticated decision would be for him or her to choose the lottery. In fact, when we look at the average lottery take-up rate, we see that 75 percent of the subjects chose the risky option.

6.1 Choices and Expectations

In this subsection, we pay closer attention to the prevalence of probabilistic unsophisticated behavior in our sample. Figure 5 presents histograms for the expected probability of the best lottery outcome, $E[p_{\bar{x}}]$, and for the MAP reported by the participants. We see that despite similar averages, the individual expectations have a more concentrated distribution than the individual reservation probabilities. One can see, for example, that 2.03 percent of the subjects would choose the lottery even if the probability of the best outcome were between 0 and 10 percent. This latter pattern of choice happens even if no one actually expected the probabilities to be so low. If this particular subgroup of individuals were indeed probabilistically sophisticated (according to our reductionist criteria), they should have opted for the risky option when given a chance.

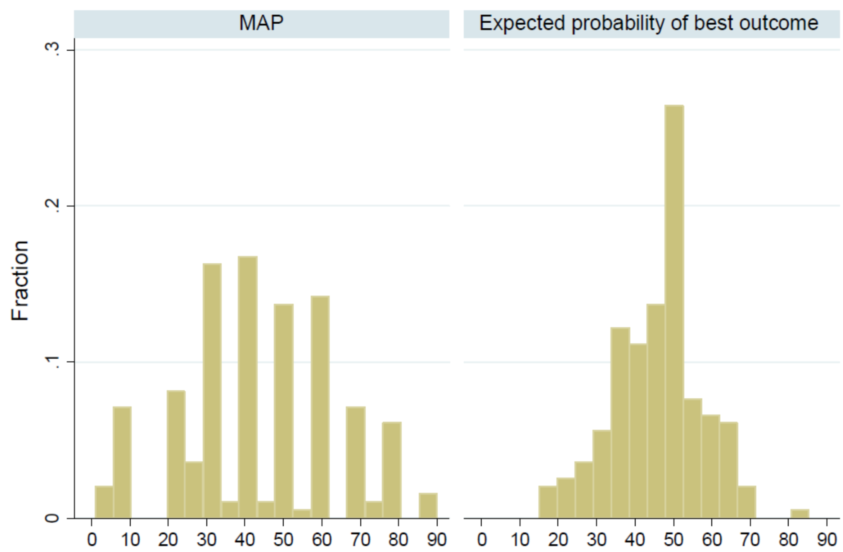

Figure 5: Histogram for MAP and Expected Probability of Lottery Option Highest Outcome

Further exploring comparisons such as the one posed above, Figure 6 examines the distribution of the difference between the MAP and the expected probabilities in the groups that chose either the safe or the risky option. Despite the overall agreement with the generally consistent level of sophistication, because more individuals who select the safe (risky) option indeed have MAPs that are larger (smaller) than their expected probability, this agreement is far from complete. In fact, 30.61 percent of the individuals who select the safe option do so despite having MAPs below their revealed expectations, whereas 43.25 percent of those taking the risky option have MAPs above their expectations.

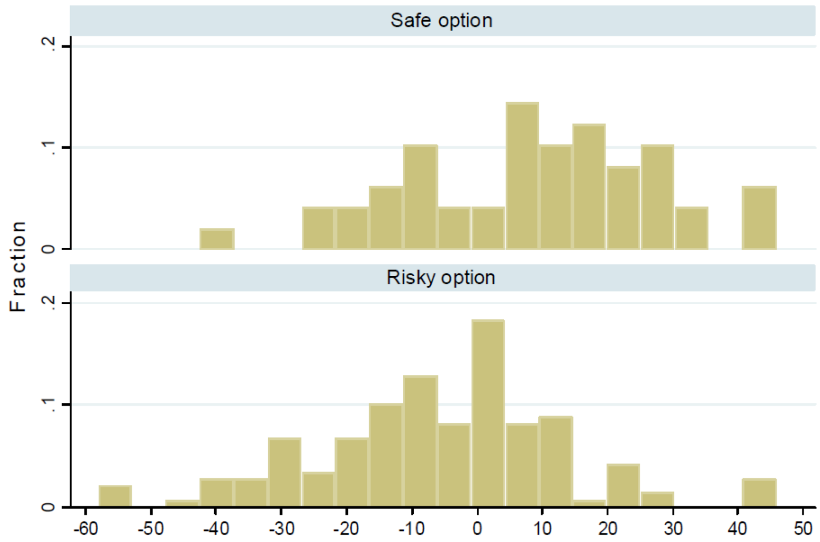

Figure 6: MAP vs. Expected Probability of Lottery Option Highest Outcome, by choice

We summarize these general findings in Table 1. In the table, we present the proportion of individuals in our sample who were found to be unsophisticated according to the three criteria proposed above. On average, 35 percent of the subjects are unsophisticated according to our weak (reductionist) definition, 12.2 percent are considered unsophisticated if we employ our hypothesis-testing criteria, and 10.7 percent are unsophisticated according to our most extreme definition of irrational behavior. The latter are people who pick an option that is only better for them outside of what they themselves consider to be the realm of possibility for the lottery outcome.

	Whole Sample Mean (SE)	CRT=3 Sample Mean (SE)	CRT<3 Sample Mean (SE)
Criterium 1 - subjects are unsophisticated according to our weak (reductionist) definition	0,350 (0,034)	0,253 (0,050)	0,415 (0,046)
Criterium 2 - considered unsophisticated if we employ our hypothesis-testing criterium	0,102 (0,022)	0,035 (0,022)	0,144 (0,033)
Criterium 3 -people who pick an option that is only better for them outside of what they themselves consider to be the realm of possibility for the lottery outcome	0,076 (0,019)	0,025 (0,018)	0,110 (0,029)

Note: Sample is 197 subjects in whole sample, and 79 subjects with CRT = 3. Robust standard errors computed using jackknife method.

Figure 7: Violations of Probabilistic Sophistication – Sample Proportions by Violation Criteria

6.2 Cognitive Ability and Probabilistic Sophistication

We now conduct a more formal analysis of the relationship between the instances of unsophisticated behavior that we observed in the data and the individual characteristics of the study participants. We begin with descriptive evidence of the distribution of fluid and crystallized intelligence (as measured by our cognitive ability tests) in the subgroups of individuals categorized as sophisticated or unsophisticated based on the three criteria suggested in the text. As shown in Figures 8 and 9, there is little doubt of the substantial differences between the groups. According to both measures, sophisticated individuals are more likely to have cognitive ability distributions that are "shifted to the right" relative to those of their unsophisticated counterparts.

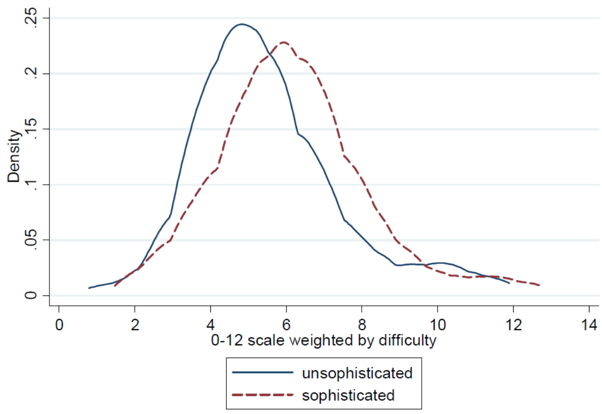

Figure 8: Density Distribution of Cognitive Ability Scores

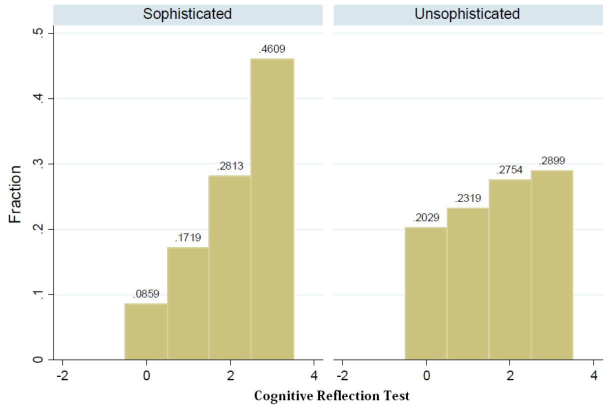

Figure 9: Histogram of Cognitive Ability Scores

We present these results in a more structured way in the Table 1 below. The panels A-C in the table present the results of the estimations of simple linear probability models that examine the influence of cognitive ability on the probability of unsophisticated behavior.[9] Each column presents the estimated quantities based on alternative specifications. From column [1] to column [10], progressively more controls are included. Our findings indicate that in all specifications, cognitive abilities can be considered powerful explanatory variables. Panel A indicates that both fluid and crystallized intelligence can be considered to be equally relevant to sophistication in its simplest format. This observation holds even after demographic, training and session characteristics are included in the regressions.

[9]Logit and probit models yield the same qualitative results and are available upon request.

Table 1: Probabilistic Sophistication and Cognitive Ability – Regression Estimates

	Without controls					With controls				
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]
Panel A: Reducionist Criterium (mean value of dependet variable = 0.350254)										
Fluid intelligence	-0,069			-0,052	-0,058	-0,057			-0,041	-0,047
	(0,038)			(0,036)	(0,037)	(0,037)			(0,036)	(0,037)
CRT		-0,101		-0,092			-0,092		-0,084	
		(0,040)		(0,040)			(0,041)		(0,042)	
CRT=3			-0,162		-0,144			-0,135		-0,119
			(0,073)		(0,073)			(0,068)		(0,069)
1{Has taken Stats course}						-0,215	-0,191	-0,203	-0,187	-0,198
						(0,099)	(0,099)	(0,097)	(0,099)	(0,096)
1{Economics}						-0,030	0,015	0,017	0,005	0,004
						(0,102)	(0,076)	(0,085)	(0,082)	(0,091)
Panel B: Hypothesis-Test Like Criterium (mean value of dependet variable = 0.101523)										
Fluid intelligence	-0,049			-0,043	-0,041	-0,042			-0,038	-0,036
	(0,018)			(0,017)	(0,017)	(0,019)			(0,018)	(0,018)
CRT		-0,039		-0,031			-0,027		-0,019	
		(0,021)		(0,020)			(0,020)		(0,018)	
CRT=3			-0,106		-0,093			-0,084		-0,071
			(0,043)		(0,041)			(0,037)		(0,034)
1{Has taken Stats course}						-0,068	-0,065	-0,062	-0,061	-0,058
						(0,046)	(0,046)	(0,048)	(0,046)	(0,048)
1{Economics}						-0,112	-0,095	-0,082	-0,104	-0,092
						(0,065)	(0,063)	(0,061)	(0,065)	(0,063)
Panel C: Disbelief in Beliefs Criterium (mean value of dependet variable = 0.076142)										
Fluid intelligence	-0,039			-0,034	-0,033	-0,029			-0,026	-0,024
	(0,015)			(0,015)	(0,014)	(0,015)			(0,015)	(0,015)
CRT		-0,031		-0,024			-0,020		-0,015	
		(0,018)		(0,018)			(0,018)		(0,018)	
CRT=3			-0,085		-0,074			-0,064		-0,055
			(0,042)		(0,040)			(0,037)		(0,036)
1{Has taken Stats course}						-0,083	-0,081	-0,078	-0,078	-0,075
						(0,043)	(0,045)	(0,046)	(0,045)	(0,046)
1{Economics}						-0,082	-0,069	-0,059	-0,076	-0,066
						(0,059)	(0,059)	(0,055)	(0,060)	(0,057)

Note: Sample is 197 subjects in whole sample. Robust standard errors are clustered at experimental session level.

Nonetheless, Panels B and C deliver a slightly modified conclusion. Despite the joint significance of the results of our cognitive tests, most of the action is a function of the CRT scores. This finding suggests that extreme cases of unsophistication originate from a lack of crystallized intelligence or from excessively impulsive decision making. Interestingly, using stricter criteria for probabilistic unsophistication, we find that training plays a relevant role and that exposure to economics and advanced science courses are particularly important. Taken together, our results confirm the association between cognitive ability and probabilistic sophistication in decision making under uncertainty.

7 How is probabilistic sophistication bounded by cognitive ability? A simple framework

Our results are consistent with those of previous work that has indicated that other aspects of individual behavior, such as risk and time preferences and the incidence of behavioral bias, are related to cognitive ability (Oechssler *et al.*, 2009, Benjamin *et al.*, 2012). However, it is natural to wonder through what mechanism cognitive ability affects probabilistic sophistication and, more generally, other forms of "rational" behavior.

Although a growing number of empirical studies have documented the association between cognitive ability and behavioral aspects that are of economic importance, little is known yet about why that association exists. In this section, we present a simple abstract model of a decision-making procedure that will help us to determine why some individuals fail to meet the prediction of normative decision theory. Our theoretical analysis should be viewed as a first step toward a more comprehensive framework for analyzing how cognitive ability is connected to deviations from theories of rational choice, even in environments that provide adequate opportunities and incentives for thoughtful deliberation.

Our model is built upon the notion that a decision maker considers the following two factors when making decisions: the benefits of adopting a more reflective cognitive mode (as opposed to an intuitive one) and the cognitive costs of doing so. We thus model the cognitive mode that agents choose in solving a generic decision problem as a simple benefit-cost problem. This model builds on earlier work by Sloman (1998), Stanovich & West (1996, 2000), and Kahneman (2003, 2011) that have proposed that there are two generic modes of cognitive function; one mode that, loosely speaking, is based on intuition and the automatic experience of perception; and another that is based on effortful and deliberate reasoning. We assume that using a reflective mode is more costly but is a necessary (and not sufficient) condition for making decisions that are consistent with normative theories of rational decision. Building on this approach could ultimately shed light on a wider range of aspects of economic behavior. In a growing body of literature, these aspects of behavior have been shown to be associated with cognitive ability. They include time preferences, wealth accumulation and risk-taking behavior, to name just a few (see references given in the Introduction).

Preliminaries. Each individual is endowed with a given cognitive ability $\theta \in \Theta$. Nature will map each θ onto a cognitive cost function $c_\theta \in C$. The cognitive function $c_\theta : E \to M$ ($M \subset R_+$) is a strictly increasing and smooth mapping function through which a given cognitive effort e leads to a cost. We assume that C is a family of smooth functions of this type, each of which satisfies the following:

(i) Thinking is costly: $c_\theta(0) = 0$, and $c_\theta(e) > 0 \;\; \forall\, e > 0$;

(ii) $c_\theta(e)$ is twice continuously differentiable $\forall e$;

(iii) The harder one thinks, the more it costs: $\frac{dc_\theta}{de} > 0$ and $\frac{d^2 c_\theta}{de^2} > 0$;

(iv) Heterogenous cognitive cost condition: for all $e \in E$, if $\theta < \hat{\theta}$ then $c_\theta(e)/c_{\hat{\theta}}(e)$ is increasing in e.

We assume that there are two modes of thought: one based on intuition (I) and one based on deliberate reasoning (R). Intuitive judgments are based on heuristic processes that solve problems based on answers that come more readily to mind. We assume that intuitive judgments are effortless because they are perceptual and automatic, whereas the reasoning-based mode is costly because it involves analytic procedures that are more serial and computationally demanding.[10]

To simplify the exposition, we assume that the benefit of engaging in an effortful and reflective cognitive process is described by $b : \mathbb{R}_+ \to \mathbb{R}$, a smooth, increasing, and strictly concave function for

[10]For a comprehensive characterization of these two modes of thought, see, e.g., Kahneman (2003).

any cognitive effort $e \in [0, \bar{e}]$, a cognitive effort feasibility interval. Hence, from the decision maker's point of view, there are two types of outcomes for each decision problem: $B_I = b(0)$ if the decision maker makes an intuitive judgment and $B_R(e) = b(e) - c_\theta(e)$ if he undergoes a reflective process in which he exerts cognitive effort e.

Choice of effort. In decision making, individuals choose the level of cognitive effort that maximizes the net benefit of that cognitive effort, taking into account a simple incentive compatibility constraint: namely, that the cognitive cost of making an effort $e > 0$, thus overriding the intuitive cognitive mode, must be at least as great as the expected benefit of such an effort.

$$\arg\max_{e \in \{0, \bar{e}\}} \left[\int_0^{\bar{e}} b(e) - c_\theta(e) de \right] = e^*$$

Cognitive mode. The act of choosing a cognitive mode in which to make a decision follows straightforwardly from the selection of cognitive effort: the "intuitive mode" of thought is chosen if the optimal cognitive effort, e^*, is zero, and the "reflective mode" is chosen if $e^* > 0$. Because intuitive decisions are effortless, easily accessible answers produced by intuitive thinking will control the decision-making process unless the answers are overridden by effortful and reasoning-based thought, which occurs only when the expected benefit of thinking in this manner exceeds the cost.

This model indicates how cognitive ability may restrict the ability of some individuals to meet the choice requirements of normative models. Figure 10 below illustrates this connection. The figure shows an instance in which three individuals who differ in their cognitive abilities $(\theta_1, \theta_2, \theta_3)$ make different choices regarding their optimal levels of cognitive effort, each of which has different implications for both their mode of thought and any departures from normative behavior.

Let us make the simplifying assumption that e^D is the minimum cognitive effort that is necessary to meet the requirements of a choice decision that is consistent with "rationality" as defined by some normative decision model. Recall that it is assumed that the cognitive cost functions shift downward and to the right as cognitive ability increases. Hence, as $\theta_1 > \theta_2 > \theta_3$, cognitive effort is less costly to individual 1 than to individuals 2 and 3. (e_1^*, e_2^*, e_3^*) is the optimal effort profile of these individuals given the expected benefit of cognitive effort $B(e)$. Note that individual 3 does not possess the necessary "cognitive machinery" to expend cost-effective cognitive effort and thus chooses the intuitive mode, which entails automatic decision making with no cognitive effort. Although individuals 2 and 3 both make decisions in a reflective mode, only individual 1 will expend a level of cognitive effort that is consistent with the "rationality" requirements of the normative model.

Although some key features of this model are quite general and abstract, the rationale governing the model is simple: as long as individuals differ sufficiently in their cognitive abilities, some individuals will depart from normative models, whereas others will not. This simple model can be used to organize the results of our experiment, explaining the departures from the behavior predicted by normative models as a function of differential cognitive ability and, consequently, of the cost-effectiveness of engaging in effortful and reflective decision making. This finding shows how intuition and reasoning guide judgment and decision making.

Figure 10: How Cognitive Ability Constrains "Rationality"

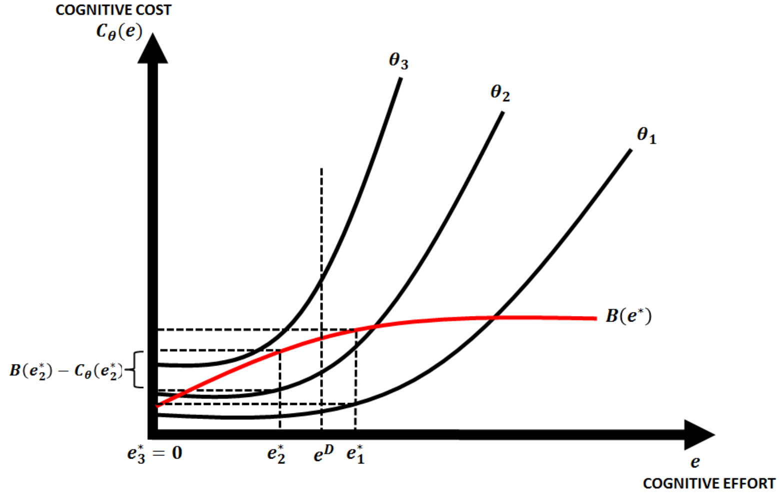

8 Concluding Remarks

A decision maker is probabilistically sophisticated if his preferences over subjectively uncertain bets are consistent with his preferences over objectively uncertain versions of those bets constructed from his probabilistic beliefs about the states that determine the outcomes of the uncertain bets. This finding is a crucial assumption in a family of models of decision making under uncertainty that assume that subjective beliefs can substitute for objective probability in the evaluation of a bet. Although this assumption has been empirically challenged by several experiments, these experiments have a remarkably similar feature: beliefs are shown to be inconsistent across sources of uncertainty that are not symmetric, that is, they involve known and unknown probabilities. In the present paper, we design and implement an experiment that demonstrates that individuals are actually probabilistically sophisticated when the source of uncertainty that they face remains *symmetric* (information-wise). This experiment was conducted in a simple decision setting in which a choice is made between a sure outcome and an uncertain bet whose actual outcome depends on states of nature for which no explicit probability of occurrence is given. We then elicited the participants' preferences over the objective lottery version of this choice and determined their probabilistic beliefs about the states that determine the outcomes of the uncertain bet. We find that a fairly large proportion of the subjects exhibit choice behavior that is consistent with probabilistic sophistication. This finding suggests that the difficulties associated with using beliefs indicated by preferences as substitutes for objective probability are not as general as has been suggested by Ellsberg-like experiments.

We also find robust evidence that cognitive ability can mitigate the violations of the requirement of probabilistic sophistication. Individuals with relatively high levels of cognitive ability are significantly less likely to make a choice under uncertainty that is inconsistent with their own probabilistic beliefs. In fact, when the consistency condition imposed on choices and beliefs is strengthened so that only cases in which the distribution of probabilistic beliefs is substantially inconsistent with the choices made are deemed violations of probabilistic sophistication, the effect of cognitive ability on the likelihood of probabilistically sophisticated behavior is even stronger.

Many economists have argued that the violations of theories of decision making that have been observed in laboratory experiments should vanish or decrease if adequate conditions for decision are secured, such as, for example, opportunities for learning and proper incentives for deliberation. Our results suggest, however, that deviations from the behavior predicted by theories of rational decision cannot be explained by purely *environmental* conditions. Deviations from probabilistic sophistication, a key feature of rational choice, are non-existent or are severely reduced among individuals with better cognitive endowments, regardless of the conditions under which the decisions were made.

Instead, differences exist in the way relatively high- and low cognitive-ability individuals compute choices, which makes the former group more likely than the latter to behave in a manner that is consistent with the predictions of theories of rational choice for that domain. In a simple cost-benefit analysis, we propose that cognitive ability generates the cognitive cost of decisions and that cognitive-based decisions are either too costly for some individuals or are not worth making for those individuals given the perceived benefits. In other words, the success that individuals with relatively high cognitive ability experience in making choices that are consistent with theories of rational choice reflects the cost-effectiveness of their cognitive machinery.

Overall, our results contribute to a recent and growing body of economic literature that demonstrates how cognitive ability can influence the behavior of interest, from risk and time preferences to behavior in strategic games. However, before we can be certain that decisions that deviate from rational choice are made because of the poor cost-effectiveness of cognitive effort, we must better determine why and under what circumstances cognitive ability shapes economic behavior. We view the present article as an initial step toward establishing a research agenda that includes more elaborate experiments and models that directly address the causal link between cognitive ability, probabilistic sophistication in particular, and the other economic aspects of individual decision making in general.

A Proof of Prediction

Because claims (i) and (ii) are a symmetric reversal of each other, we prove claim (i) here only. We must demonstrate that the belief that the act g will yield the "good" prize x with probability $\mu(s) > MAP$ should prompt a probabilistically sophisticated individual to prefer g over h_z.

We begin with the fact that $g : I \mapsto L$: that is, the act g maps each $i \in I = \{0, 1, 2, ..., 10\}$ onto a lottery $\ell_i \in L = \{\ell_i(x, \frac{i}{10}, y) : i \in I\}$. The definition of $\mu(.)$ immediately implies that $\mu(i)$ maps g onto a compound lottery $L^*\Big(\ell_0, \mu(0), \ell_1, \mu(1), \ldots, \ell_{10}, \mu(10)\Big)$. Based on the principle of reduction for compound lotteries, a lottery ℓ^* whose outcomes are probability distributions over x and y can be reduced to a binary lottery $\ell \in L$ via the multiplication of the objective probabilities by the subjective ones. Thus $L^*\Big(\ell_0, \mu(0), \ell_1, \mu(1), \ldots, \ell_{10}, \mu(10)\Big) = \ell^*(x, p^*, y)$ where $p^* = \sum_{i \in I} \mu(i) \dfrac{i}{10}$. By first-order stochastic dominance, if

$\ell(x, q, y) \succ \ell(z, 1) \quad \forall \quad q > MAP$, then $\ell^*(x, p^*, y) \succ \ell(z, 1) \quad \forall \quad p^* > MAP$.

If two acts m and n, through $\mu(.)$, map onto lotteries ℓ_m and ℓ_n and $\ell_m \succ \ell_n$, probabilistic sophistication dictates that $m \succ n$. Therefore, if $\ell^*(x, p^*, y) \succ \ell(z, 1)$, then $g \succ h_z$ whenever $p^* > MAP$. This completes the proof.

References

Abdellaoui, Mohammed, Baillon, Aurelien, Placido, Laetitia, & Wakker, Peter P. 2011. The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation. *American Economic Review*, **101**(2), 695–723.

Anscombe, F., & Aumann, R. J. 1963. A Definition of Subjective Probability. *Annals of Mathematical Statistics*, **34**(1), 199–205.

Benjamin, Daniel J., Brown, Sebastian A., & Shapiro, Jesse M. 2012. Who is "Behavioral"? Cognitive Ability and Anomalous Preferences. *Journal of the European Economic Association*, **Forthcoming**.

Blanco, Mariana, Engelmann, Dirk, Koch, Alexander, & Normann, Hans-Theo. 2010. Belief elicitation in experiments: is there a hedging problem? *Experimental Economics*, **13**(4), 412–438.

Bohnet, Iris, Greig, Fiona, Herrmann, Benedikt, & Zeckhauser, Richard. 2008. Betrayal Aversion: Evidence from Brazil, China, Oman, Switzerland, Turkey, and the United States. *American Economic Review*, **98**(1), 294–310.

Camerer, Colin, & Weber, Martin. 1992. Recent Developments in Modeling Preferences: Uncertainty and Ambiguity. *Journal of Risk and Uncertainty*, **5**(4), 325–70.

Carrol, John Bissell. 1993. *Human Cognitive Abilities: A Survey of Factor-Analytic Studies.* New York: Cambridge University Press.

Cattell, R. B. 1943. The measurement of adult intelligence. *Psychological Bulletin*, **40**, 153–193.

Cohen, M., Jaffray, J. Y., & Said, T. 1985. Individual behavior under risk and under uncertainty: An experimental study. *Theory and Decision*, **18**, 203–228.

de Finetti, Bruno. 1937. La prvision:ses lois logiques, ses sources subjectives. *Ann. Inst. H. Poincarr*, **7**, 1–68.

Ellsberg, Daniel. 1961. Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, **75**(4), 643–669.

Fishburn, Peter C. 1969. A General Theory of Subjective Probabilities and Expected Utilities. *The Annals of Mathematical Statistics*, **40**(4), 1419–1429.

Fox, Craig R, & Tversky, Amos. 1995. Ambiguity Aversion and Comparative Ignorance. *The Quarterly Journal of Economics*, **110**(3), 585–603.

Frederick, Shane. 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, **19**(4), 25–42.

Gilboa, Itzhak, & Schmeidler, David. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, **18**(2), 141–153.

Grant, Simon. 1995. Subjective Probability without Monotonicity: Or How Machina's Mom May Also Be Probabilistically Sophisticated. *Econometrica*, **63**(1), 159–89.

Halevy, Yoram. 2007. Ellsberg Revisited: An Experimental Study. *Econometrica*, **75**(2), 503–536.

Heath, Chip, & Tversky, Amos. 1991. Preference and Belief: Ambiguity and Competence in Choice under Uncertainty. *Journal of Risk and Uncertainty*, **4**(1), 5–28.

Horn, J. L., & Cattell, R. B. 1967. Age differences in fluid and crystallized intelligence. *Acta Psychologica*, **26**, 107–129.

Kahn, Barbara E, & Sarin, Rakesh K. 1988. Modeling Ambiguity in Decisions under Uncertainty. *Journal of Consumer Research*, **15**(2), 265–72.

Kahneman, Daniel. 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *American Economic Review*, **93**(5), 1449–1475.

Kahneman, Daniel. 2011. *Thinking, fast and slow.* Farrar, Straus and Giroux.

Karni, Edi. 2009. A Mechanism for Eliciting Probabilities. *Econometrica*, **77**(2), 603–606.

MACHINA, MARK J, & SCHMEIDLER, DAVID. 1992. A More Robust Definition of Subjective Probability. *Econometrica*, **60**(4), 745–80.

MARINACCI, MASSIMO. 2002. Probabilistic Sophistication and Multiple Priors. *Econometrica*, **70**(2), 755–764.

OECHSSLER, JRG, ROIDER, ANDREAS, & SCHMITZ, PATRICK W. 2009. Cognitive abilities and behavioral biases. *Journal of Economic Behavior & Organization*, **72**(1), 147–152.

OFFERMAN, THEO, SONNEMANS, JOEP, KUILEN, GIJS VAN DE, & WAKKER, PETER P. 2009. A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes-super-. *Review of Economic Studies*, **76**(4), 1461–1489.

PRATT, JOHN W., RAIFFA, HOWARD, & SCHLAIFER, ROBERT. 1964. The Foundations of Decision Under Uncertainty: An Elementary Exposition. *Journal of the American Statistical Association*, **59**(306), 353–375.

RAMSEY, FRANK P. 1926. Truth and Probability. *Chap. 7, pages 156–198 of:* BRAITHWAITE, R. B. (ed), *The Foundations of Mathematics and other Logical Essays*. McMaster University Archive for the History of Economic Thought.

RAMSEY, FRANK P. 1931. *The Foundations of Mathematical and Other Logical Essays*. Routledge and Keagan Paul. Chap. Truth and Probability, pages 156–198.

RAVEN, J. C., J. H. COURT. 1998. *Manual for Raven's Progressive Matrices and Vocabulary Scales: Section 1*. USA: Harcourt.

SAVAGE, L. 1954. *The Foundations of Statistics*. New York: Wiley.

SCHMEIDLER, DAVID. 1989. Subjective Probability and Expected Utility without Additivity. *Econometrica*, **57**(3), 571–87.

SLOMAN, STEVEN A. 1998. The empirical case for two systems of reasoning. *Journal of Experimetnal Psychology: General*, **127**, 161–188.

STANOVICH, KEITH E., & WEST, RICHARD F. 1996. Individual differences in rational thought. *Journal of Experimental*, **119**, 3–22.

STANOVICH, KEITH E., & WEST, RICHARD F. 2000. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, **23**(05), 645–665.

STRZALECKI, TOMASZ. 2011. Probabilistic Sophistication and Variational Preferences. *Journal of Economic Theory*, **146**, 2117–1194.

TRAUTMANN, STEFAN T., & VAN DE KUILEN, GIJS. 2010. *Belief Elicitation: A Horse Race among Truth Serums*. Working Paper. Tilburg University.