NBER WORKING PAPER SERIES

COGNITIVE UNCERTAINTY

Benjamin Enke
Thomas Graeber

Working Paper 26518
http://www.nber.org/papers/w26518

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2019, Revised January 2021

The experiments in this paper were pre-registered in the AEA RCT registry as trial AEARCTR-0004493. Graeber thanks the Sloan Foundation for Post-Doctoral Funding and Enke the Foundations of Human Behavior Initiative for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Cognitive Uncertainty
Benjamin Enke and Thomas Graeber
NBER Working Paper No. 26518
November 2019, Revised January 2021
JEL No. D01,D03

ABSTRACT

Because many economic decisions are difficult, people may exhibit cognitive uncertainty: subjective uncertainty about what the optimal action is. This paper shows that cognitive uncertainty predicts economic beliefs and actions, and that it provides a unifying lens for understanding behavioral anomalies in how people think about probabilities. The main idea is that when people are cognitively uncertain, they act as if they compress objective probabilities towards a cognitive default that is given by an ignorance prior. By experimentally measuring and exogenously manipulating cognitive uncertainty in different decision contexts, our analysis brings together and partially explains a large set of empirical regularities in choice under risk, choice under ambiguity, belief updating, and survey forecasts of economic variables. These include the probability weighting function, the fourfold pattern of risk attitudes, ambiguity-insensitivity, base rate insensitivity, conservatism, sample proportion effects, and predictable overoptimism and -pessimism in economic forecasts. Because people's reported cognitive uncertainty systematically varies as a function of the objective probabilities in a decision problem, our framework also sheds light on the pronounced inverse S-shaped response patterns that pervade different literatures in behavioral economics.

Benjamin Enke
Department of Economics
Harvard University
Littauer M008
Cambridge, MA 02138
and NBER
enke@fas.harvard.edu

Thomas Graeber
Harvard University
Department of Economics
graeber@fas.harvard.edu

1 Introduction

In many contexts of economic interest, decision-making under uncertainty is difficult. In belief formation, people may not know Bayes' rule, succumb to computational errors, or struggle to retrieve and integrate all relevant information. In choice under risk, people may not know their true preferences, or fail at adequately combining probabilities and utils. These issues, and potentially many more, may introduce cognitive noise, which we use as a catch-all term for unsystematic errors that arise from cognitive imperfections in the process of optimization.

Our basic premise is that people are often aware of their own cognitive noise, which induces *cognitive uncertainty*: subjective uncertainty about what the optimal action or solution to a decision problem is. For example, people may think that they do not really know their own certainty equivalent of a lottery; they may have a nagging feeling that they do not remember what their prior information is; or they may worry that they do not know how to compute rational beliefs in light of new information. Indeed, recent work in psychology and neuroscience on decision confidence suggests that people often have a sense of how "good" their decision is.

The objective of this paper is to document empirically that cognitive uncertainty predicts economic beliefs and actions, and that it provides a unifying lens for understanding well-known empirical regularities in how people think about probabilities. The key idea is that *noise and bias are linked*: when people are cognitively noisy, they revert more to a cognitive default, which introduces systematic bias. According to this argument, the "second moment" of people's decisions (their confidence in what they are doing) is relevant for understanding the "first moment:" which decisions they take in the first place.

Figure 1 illustrates the set of well-established empirical regularities that we focus on. All four functions are estimated from experimental data and share in common a characteristic inverse S-shape of subjective with respect to objective probabilities. First, panel A depicts the well-known probability weighting function in choice under risk that goes back to Tversky and Kahneman (1992). It illustrates how experimental subjects implicitly treat objective probabilities in choosing between different monetary gambles. Second, panel B shows an "ambiguity weighting function" that depicts the emerging consensus that people are ambiguity averse for likely gains, yet ambiguity seeking for unlikely gains. This reflects a compression effect that is labeled "a-insensitivity" in the literature (Trautmann and Van De Kuilen, 2015). Third, in panel C, we illustrate a perhaps less well-known stylized fact, which is an inverse S-shaped relationship between participants' posterior beliefs and the Bayesian posterior in canonical "balls-and-urns" belief updating tasks. Finally, panel D of Figure 1 shows the relationship between objectively

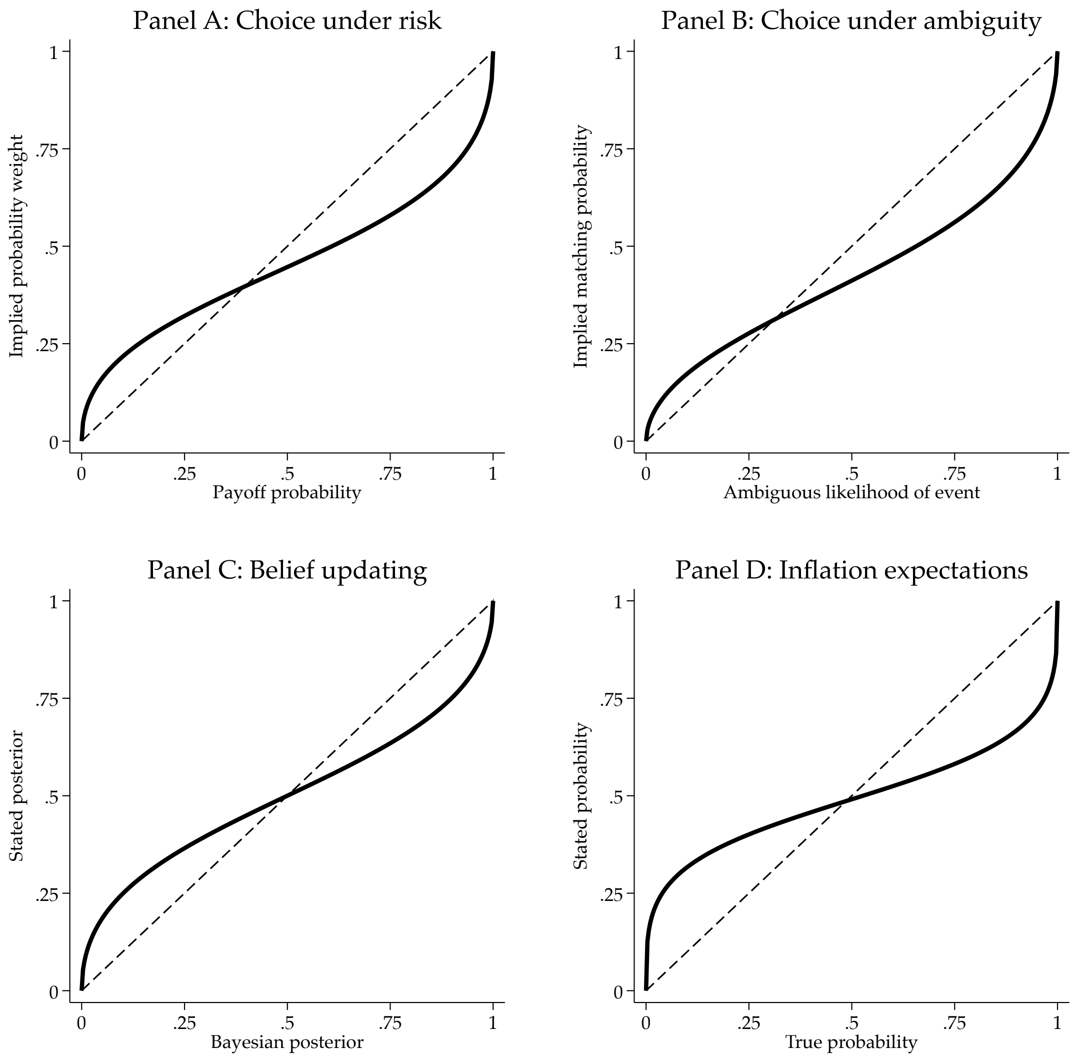

Figure 1: "Weighting functions" in choices and beliefs. Panel A depicts a probability weighting function, estimated from the data described in Section 3. Panel B illustrates an "ambiguity weighting function," where the x-axis represents the ambiguous likelihood of an event and the y-axis the matching probability (adapted from Li et al., 2019). Panel C visualizes the relationship between stated beliefs and Bayesian posteriors in binary-state balls-and-urns belief updating experiments, constructed from the data described in Section 4. Finally, panel D depicts the relationship between stated subjective probabilities in a survey on inflation expectations and objective probabilities, as described in Section 5.

correct probabilities and respondents' probabilistic estimates in subjective expectations surveys about, e.g., stock market returns, inflation rates, or the shape of the income distribution. Here, again, people's beliefs are compressed towards 50:50 (Fischhoff and Bruine De Bruin, 1999). Why do these four functions, drawn from different decision contexts and experimental paradigms, look so strikingly similar?

To address this question, we present a series of experiments in which we link choices and beliefs to an empirical measure of cognitive uncertainty. While we view our primary contribution as experimental, we structure our analysis through a theoretical framework that follows an emerging class of Bayesian noisy cognition models, in particular Gabaix (2019) and Khaw et al. (2017). We take a broad interpretation of these models as capturing cognitive noise that primarily results from high-level reasoning in optimization rather than perceptual imperfections alone. In the model, people exhibit cognitive noise in translating probabilistic information into an optimal response. This cognitive noise induces people to shrink objective probabilities towards a prior, or cognitive default. While the cognitive default in general likely depends on a multitude of factors, we assume (and experimentally confirm) that in unfamiliar environments it is influenced by an ignorance prior, which assigns equal probability to all states of the world.

Given this setup, we formally define an empirically measurable notion of cognitive uncertainty as subjective uncertainty about the optimal action. Our model endogenizes the well-known neo-additive weighting function: a decision-maker's action can be expressed as a linear combination of objective probabilities and the cognitive default, where the relative weights are predicted by the magnitude of cognitive uncertainty. Endogenizing the weighting function clarifies that (i) we expect it to accurately predict behavior not just in choice under risk but also in belief formation and (ii) that the slope and elevation parameters of this function will depend on the magnitude of cognitive noise and the location of the cognitive default.

While our framework is deliberately kept stylized and does not feature the richness of domain-specific models, it allows us to transparently illustrate the logic of cognitive compression and to formally define cognitive uncertainty. Substantively, this theoretical framework makes six predictions: (a) people state positive cognitive uncertainty; (b) subjective probabilities implied by average actions are biased towards the cognitive default (50:50 in binary state spaces), which leads to a compression effect; (c) correlationally, individuals with higher cognitive uncertainty compress probabilities more towards 50:50; (d) an exogenous increase in cognitive uncertainty generates more compression in subjective probabilities; (e) an exogenous decrease in the location of the cognitive default shifts subjective probabilities downwards across the entire probability range; and (f) individuals with higher cognitive uncertainty react more strongly to exogenous variation in the location of the cognitive default.

To test these predictions, we implement a series of pre-registered experiments with a total of $N = 2,800$ participants on Amazon Mechanical Turk (AMT). Like the motivating examples, our experiments cover the domains of choice under risk and ambiguity, balls-and-urns belief updating tasks, and survey expectations about economic variables. We always work with a two-step procedure, whereby we first elicit experimental actions in a standard fashion and then measure cognitive uncertainty about these actions.

In choice under risk, we first elicit participants' certainty equivalents for two-outcome gambles such as "Get \$20 with probability 75%; get \$0 with probability 25%" in a standard price list format. Then, we measure cognitive uncertainty as participants' subjectively perceived uncertainty about the optimality of their own action. We ask participants how certain they are that to them the lottery is worth exactly the same as their revealed switching interval. To answer this question, participants use a slider to calibrate the statement "I am certain that the lottery is worth betwen x and y to me." If a subject moves the slider to the very right, x and y collapse to the switching interval in the price list. The further a subject moves the slider to the left, the wider the range of cognitive uncertainty becomes. This measure of cognitive uncertainty (i) directly reflects subjects' own assessment of uncertainty and (ii) is quantitative in nature. We discuss why for our purposes this measure is conceptually preferable to alternative measures such as the extent of across-task inconsistency.

In contrast to the predictions of rational or behavioral models without cognitive noise, our data show that about 50% of the time, subjects exhibit cognitive uncertainty that is strictly wider than the switching interval of \$1. Such cognitive uncertainty is strongly correlated with the magnitude of likelihood insensitivity in probability weighting. This implies that, as predicted by our framework, cognitive uncertainty is positively correlated with risk taking for low probability gains and high probability losses, yet negatively correlated with risk taking for high probability gains and low probability losses. Cognitive uncertainty is thus correlated with a more pronounced "fourfold pattern of risk attitudes." While our analyses explicitly embrace across-subject heterogeneity in cognitive uncertainty, we further show that measured cognitive uncertainty even predicts the magnitude of likelihood insensitivity across decisions within the same subject.

To exogenously increase cognitive uncertainty, we introduce compound and ambiguous lotteries. To illustrate, a compound lottery is a lottery that pays a non-zero amount with probability $p \sim U[0, 0.2]$. Similarly, an ambiguous lottery is a lottery that pays a non-zero amount with unknown probability $p \in [0, 0.2]$. We show that compound and ambiguous lotteries indeed induce substantially higher cognitive uncertainty than the corresponding reduced lotteries. Our model predicts that this increase in cognitive uncertainty translates into a more compressed and thus more insensitive weighting function. Our experimental results support this hypothesis: the observed likelihood insensitivity is substantially more pronounced with ambiguous or compound lotteries. As a result, subjects act as if they are "compound risk seeking" and "ambiguity seeking" for low probability gains and high probability losses.

In a final step of the analysis of choice under risk, we exogenously manipulate the location of the cognitive default. To this effect, we leverage our assumption that in unfamiliar environments the default is influenced by an ignorance prior. In the two-states lotteries discussed so far, this ignorance prior is given by 50:50. To manipulate the location of the cognitive default, we implement a partition manipulation and translate the two-states lotteries into ten-states lotteries, without changing the objective payoff profile. We hypothesize that this shifts the cognitive default to an ignorance prior of 10%, which should decrease the elevation of the entire probability weighting function. In our data, this treatment variation indeed shifts the estimated weighting function significantly towards zero. Moreover, we find that decisions that are associated with higher cognitive uncertainty respond more strongly to variation in the cognitive default, just like our cognitive shrinkage model predicts. These results on the default manipulation are important because they show that the compressed response functions that motivate our study are indeed generated by cognitive shrinkage towards a cognitive default, rather than an extreme version of random choice.

In a second set of experiments, we conduct conceptually analogous exercises for belief updating. Here, we implement canonical balls-and-urns updating tasks. In these experiments, a computer randomly selects one of two bags according to a known base rate, yet subjects do not observe which bag got selected. The two bags both contain 100 balls, where one bag contains $q > 50$ red and $(100 - q)$ blue balls, while the other bag contains q blue and $(100-q)$ red balls. The computer randomly draws one or more balls from the selected bag and shows these balls to the subject, who is then asked to provide a probabilistic assessment of which bag was actually drawn. Across experimental tasks, the base rate, the signal diagnosticity q and the number of random draws vary, but are always known to subjects. The standard finding in this literature is that participants' posterior beliefs are insufficiently sensitive to variation in the Bayesian posterior.

In our experiments, we again elicit cognitive uncertainty after participants have stated their posterior belief. In a conceptually very similar fashion to choice under risk and ambiguity, we ask subjects to use a slider to calibrate the statement "I am certain that the optimal guess is between x and y." We explain that the optimal Bayesian guess relies on the same information that is available to subjects. As a complementary, and financially incentivized, measure of cognitive uncertainty, we also elicit subjects' willingness-to-pay to replace their own guess with the optimal guess.

Again, in contradiction to a large class of models in which agents do not exhibit doubts about the rationality of their belief updating, in the vast majority of cases (86%), subjects indicate strictly positive cognitive uncertainty. As predicted by our model, this cognitive uncertainty is strongly correlated with compression of posterior beliefs towards 50:50. Again, these results hold both when we embrace across-subject variation in cognitive uncertainty and when we focus more narrowly on variation in cognitive uncertainty across decisions made by the same subject. Furthermore, we explain and document empirically how our account of cognitive uncertainty endogenizes many of the empirical regularities that the experimental literature on belief updating has accumulated, including "extremeness aversion," base rate insensitivity, likelihood ratio insensitivity (conservatism), and sample proportion effects. Our results suggest that these empirical regularities derive from a single cognitive principle.

Similarly to our choice under risk experiments, we complement these correlational belief updating results with exogenous manipulations of both cognitive uncertainty and the cognitive default. First, we again exogenously shift cognitive uncertainty using a compound manipulation, which substantially increases cognitive uncertainty and leads the distribution of beliefs to become substantially more compressed towards 50:50, as predicted by our framework. Second, we again manipulate the location of the cognitive default using a partition manipulation, without changing the relevant Bayesian posterior. Similarly to our choice under risk experiments, we again find that this treatment shifts the entire distribution of responses, where high cognitive uncertainty decisions respond more to the exogenous variation in the the cognitive default. Again, these results are inconsistent with random choice and show that our results are driven by a cognitive shrinkage towards a specific cognitive default, rather than to 50:50 per se.

In the third part of the paper, we study the relationship between cognitive uncertainty and survey expectations about the performance of the stock market, inflation rates, and the structure of the national income distribution. These expectations are conceptually slightly different from the laboratory choice under risk and belief updating tasks in that there is potentially information that participants do not have, while the lab tasks are self-contained. Nevertheless, these data allow us to assess whether our cognitive uncertainty measure also predicts compression in beliefs in more applied contexts. We indeed find that subjects with higher cognitive uncertainty state expectations that are more regressive towards 50:50.

All of our analyses have a structural interpretation in terms of the neo-additive weighting function. According to this functional form, a decision-maker's action can be expressed as a linear combination of objective probabilities and the cognitive default, where the relative weights – and thus the slope of the weighting function – are determined by the magnitude of cognitive noise. This linear functional form endogenously arises because the baseline version of our model assumes that cognitive noise is constant in objective probabilities. Yet, while this linear relationship explains the key pattern of compression towards 50:50, it does not capture the nonlinearity of the canonical inverse S-shaped response patterns summarized in Figure 1. Thus, we additionally present structural exercises in which we relax the assumption that cognitive noise is constant across the probability range. Indeed, in our data, measured cognitive uncertainty exhibits a pronounced hump shape in objective probabilities: across all decision domains, participants' reported cognitive uncertainty reveals that they find it easier to think about extreme probabilities than about intermediate ones. In a series of quantitative exercises, we show that this non-linearity in the cognitive difficulty of thinking through probabilities directly produces inverse S-shaped response functions that quantitatively match our data well.

In the last part of the paper, we document that participants appear to exhibit somewhat stable cognitive uncertainty "types:" stated cognitive uncertainty is highly correlated across tasks, both within and across choice domains. For example, participants with high cognitive uncertainty in choice under risk (or belief updating) also exhibit high cognitive uncertainty in survey expectations.

The paper proceeds as follows. Section 2 lays out a theoretical framework. Sections 3 to 6 present the main experiments and corresponding structural exercises. Sections 7 and 8 study the correlates of cognitive uncertainty and present robustness checks. Section 9 discusses related literature and concludes.

2 Theoretical Framework

Our stylized theoretical framework follows the Bayesian noisy cognition literature that has gained widespread popularity outside of economics, and has recently also been applied to economic decision-making. Our exposition builds on Khaw et al. (2017) and Gabaix (2019). The central assumption of this class of models is the existence of cognitive noise in decision-making, which induces the decision-maker to form an implicit update.[1] In contrast to some earlier work, we interpret this noise not necessarily as reflecting low-level perceptual imperfections, but as resulting primarily from higher-level reasoning during optimization. We show that *awareness* of such cognitive noise creates cognitive uncertainty: subjective uncertainty about what the optimal action is. To illustrate informally, suppose your prior belief that it rains tomorrow is 15%. Next, a weather forecast predicts that it will rain. You know from experience that the weather forecast is correct 80% of the time. What is your posterior belief that it will rain tomorrow? 45%? Really? Not 40%? Or perhaps 52%? To take another example, suppose you were asked to state your certainty equivalent of a 25% chance of getting \$15. You announce \$3. But is it really \$3? Or maybe \$2.50 or \$3.20? In these examples, the feeling of uncertainty about a posterior belief or a certainty equivalent reflects cognitive uncertainty.

Cognitive Noise, Shrinkage and their Interpretation. Consider a setup in which the rational action is given by $a^r = f(p) = Bp$ and a decision-maker takes an action a that potentially differs from a^r. Here, p is a probability and B a scaling parameter. We do not explicitly model the first principles that determine the rational action. Instead, by "action," we generically refer to the solution to a decision problem such as a stated

[1]See Viscusi (1989) for an early related model in the context of probability weighting.

posterior belief or a stated certainty equivalent. In a first application, $a^r = p$ reflects the Bayesian posterior in a belief updating task, while a is the subject's stated posterior belief. In a second application, p represents the payout probability of a binary lottery, a^r the subject's true certainty equivalent, and a the actually revealed certainty equivalent. In Appendix A, we generalize the exposition to rational actions $a^r = f(p)$ that are non-linear functions of p, such as when the decision-maker is risk-averse.[2]

The decision-maker's objective is to minimize the squared distance between his action and the rational action:[3]

$$\min_{a} \quad v(a,p) = \frac{1}{2}(a - Bp)^2 \,. \tag{1}$$

We assume that the process required to identify the rational action a^r is subject to cognitive noise. We model this as the agent receiving a signal $s = p + \varepsilon$ instead of having direct access to p, with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. We view this "noisy perception" formalization *as if*, in that it arises in the process of optimizing. In choice under risk, cognitive noise arises because combining probabilities, payouts and preferences into a certainty equivalent may be difficult. In belief updating, cognitive noise arises in the process of combining the available information into a posterior. Finally, in survey expectations, cognitive noise arises through the process of retrieving and assessing information from memory.

We do not intend to take a strong stance on how across-agent differences in the magnitude of cognitive noise σ_ε^2 arise. Plausibly, cognitive noise is reduced through deliberation or "mental simulation" of the problem, as in models of sequential "evidence" accumulation. According to such a perspective, lower cognitive skills, inattention, fast responses, cognitive load or a higher complexity of the decision problem could all increase cognitive noise (see the discussion in Gabaix and Laibson, 2017). Below, we study some of these factors empirically.

The agent holds a prior $p \sim \mathcal{N}(p^d, \sigma_p^2)$, where we refer to p^d as the "cognitive default." We assume normally distributed variables throughout for tractability, but acknowledge that this assumption has limited realism for the case of probabilities, which are bounded by 0 and 1. While in general the prior is likely to be be influenced by a multitude of factors, in our empirical applications we will operate under the assumption that it is influenced by an ignorance prior that assigns equal mass to all states of the world.[4] This assumption is attractive in typical experimental applications, with which

[2]For instance, for a decision-maker with Bernoulli utility function $u(x)$, the certainty equivalent of a p% chance of receiving \$1 is given by $CE = u^{-1}(p)$, so that $f(\cdot) = u^{-1}(\cdot)$. We explicitly incorporate risk aversion in our estimations in Section 6.

[3]Our results do not depend on a quadratic loss function. The decision-maker plays the posterior mean when loss is quadratic, the posterior median for the absolute loss function, and the posterior mode for the relaxed 0-1 loss, all of which are identical in our Gaussian setup.

[4]Such an ignorance prior may be related to the well-known 1/N heuristic (Benartzi and Thaler, 2001).

people have limited prior experience. In what follows, we assume that all agents have the same prior distribution but potentially differ in the magnitude of cognitive noise σ_ε^2.

Agents account for their cognitive noise by forming an implicit update about p. For a Bayesian agent, this creates a standard Gaussian signal extraction problem:

$$\mathbb{P}(p|s) \sim \mathcal{N}(\lambda s + (1-\lambda)p^d, (1-\lambda)\sigma_p^2), \tag{2}$$

with the shrinkage factor

$$\lambda = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_\varepsilon^2} \in [0,1]. \tag{3}$$

An agent with cognitive noise who is otherwise rational takes an action by solving: $\max_a \mathbb{E}\left[-\frac{1}{2}(a - Bp)^2 |s\right]$, leading to an observed action

$$\hat{a} = B(\mathbb{E}[p|s]) = B\left[\lambda p + \lambda \varepsilon + (1-\lambda)p^d\right]. \tag{4}$$

This should be compared with the rational action $a^r = Bp$. We see that the agent dampens his response to p by λ, generating shrinkage towards the cognitive default. The takeaway is that cognitive noise makes the otherwise-rational action (i) *less sensitive* to variation in p, yet (ii) *excessively sensitive* to the normatively irrelevant variation in p^d.

Cognitive Uncertainty. Awareness of cognitive noise generates subjectively perceived uncertainty about what the *truly* optimal action is. We label this *cognitive uncertainty*. The agent's cognitive uncertainty takes as given his individual draw of s, and reflects how the truly optimal action $a^o|s$ (that is unknown to the agent) subjectively varies due to his posterior uncertainty about p (equation (2)). While the agent's loss function induces him to play $\hat{a} = \mathbb{E}[a^o|s]$ (equation (4)), the underlying perceived posterior distribution of the optimal action is

$$a^o|s \sim \mathcal{N}(B\lambda s + B(1-\lambda)p^d, B^2(1-\lambda)\sigma_p^2). \tag{5}$$

Definition. *The agent's cognitive uncertainty is given by*

$$\sigma_{CU} = \sigma_{a^o|s} = |B|\sqrt{1-\lambda}\,\sigma_p = |B|\frac{\sigma_\varepsilon \sigma_p}{\sqrt{\sigma_\varepsilon^2 + \sigma_p^2}}. \tag{6}$$

Under our maintained assumption that the prior distribution is fixed across agents, all variation in σ_{CU} reflects heterogeneity in cognitive noise. Note from (6) that higher cognitive uncertainty is associated with more shrinkage to the default (lower λ).

Discussion of Setup. Our theoretical framework is evidently stylized in nature. It is not meant to reflect a structural model of decision-making across the different decision domains that we consider in this paper. For instance, when applied to choice under risk, our model does not prescribe what the agent's true certainty equivalent is. Similarly, when applied to belief formation, our model does not involve an explicit formulation of how the agent combines prior beliefs and signals into a posterior. Lastly, the quadratic loss function does not directly correspond to the financial incentives present in experiments. We do not intend to deny that modeling these elements is important when the researcher's interest is in understanding a specific decision domain in detail. Our objective, however, is to highlight *commonalities in how people process probabilities across different decision domains*. This implies that we focus on the response patterns that emerge when the decision-maker struggles with identifying the rational action, regardless of the principles that underlie this rational action in the first place. Our deliberately simple framework allows us to transparently spell out commonalities, to formalize the link between cognitive noise and the compression of actions, and to formally define the notion of cognitive uncertainty.

In terms of interpretation, in literal terms our model posits shrinkage of the "input" quantity p. However, the model also permits an equivalent interpretation of shrinkage of the action a.[5] According to this interpretation, the agent has a cognitive default in action space (e.g., a canonical posterior belief in a belief updating task) and then combines this default with a noisy signal about the rational action.

Predictions. Normalizing $B = 1$, the expectation of equation (4) corresponds to the neo-additive weighting function:

$$w(p)^{neo} := [1 - \lambda(\sigma_{CU})] \cdot p^d + \lambda(\sigma_{CU}) \cdot p = \delta + \lambda \cdot p. \tag{7}$$

As discussed by Wakker (2010), this weighting function is appealing due its simplicity and because it can be estimated through simple linear regressions. Our model motivates this functional form by endogenizing its parameters, where the slope λ and the intercept δ depend on cognitive noise.[6]

[5]Formally, using $a^r(p) = Bp$ and denoting $a^d = Bp^d$, we get $\mathbb{E}[\hat{a}] = \mathbb{E}\left[B\left(\lambda p + \lambda\varepsilon + (1-\lambda)p^d\right)\right] = \lambda a^r(p) + (1-\lambda)a^d$.

[6]It may be helpful to contrast equation (7) with the response functions in traditional random choice models. Applying these models to our context, the agent's action is given by some version of $a = Bp + \varepsilon$. Whenever the action scale is bounded (e.g., $a \in [0,1]$), these random choice models also generate mechanical "compression" towards the center of the response scale, in expectation. The key differences are that in our model (i) actions are compressed towards the cognitive default, not necessarily the center of the response scale (we empirically distinguish between these two below); and (ii) compression of average actions does not reflect bunching at the boundaries but instead the response to cognitive noise that affects the entire distribution of stated responses (including the modal action).

Prediction 1. *Higher measured cognitive uncertainty is associated with more compressed weighting functions (lower estimated λ). That is, higher cognitive uncertainty is associated with a lower sensitivity of actions and beliefs to objective probabilities.*

Prediction 2. *An exogenous increase in cognitive noise induces more compressed weighting functions (lower estimated λ).*

Prediction 3. *An exogenous decrease in the cognitive default induces weighting functions with lower elevation (lower estimated δ).*

Prediction 4. *Higher cognitive uncertainty is associated with a higher sensitivity of actions and beliefs to exogenous variation in the cognitive default.*

It is worth comparing predictions 1 and 4. Because agents with higher cognitive uncertainty place a higher weight on the cognitive default in their decisions, their actions are *less* responsive to variation in objective probabilities but *more* responsive to variation in the (normatively irrelevant) cognitive default. This clarifies that high cognitive uncertainty does not imply a generic "insensitivity to everything."

3 Choice Under Risk

3.1 Experimental Design

Our experimental designs are guided by two objectives. First, to replicate standard choice designs from the literature to make our results comparable. Second, to propose a quantitative measure of cognitive uncertainty that is readily portable across decision domains, and reasonably easy to implement. For these reasons, we work with a two-step procedure, whereby we first elicit standard actions and then cognitive uncertainty.

3.1.1 Measuring Choice Behavior

To estimate a probability weighting function, we follow a large set of previous works and implement price lists that elicit certainty equivalents for lotteries (see, e.g. Tversky and Kahneman, 1992; Bruhin et al., 2010; Bernheim and Sprenger, 2019). In treatment *Baseline Risk*, each subject completed a total of six price lists. On the left-hand side of the decision screen, a simple lottery was shown that paid y with probability p and nothing otherwise. On the right-hand side, a safe payment z was offered that increased by \$1 for each row that one proceeds down the list. As in Bruhin et al. (2010) and Bernheim and Sprenger (2019), the end points of the list were given by $z = \$0$ and $z = \$y$.

Throughout, we do not allow for multiple switching points. This facilitates a simpler elicitation of cognitive uncertainty, as discussed below. To enforce unique switching points, we implemented an auto-completion mode: if a subject chose Option A in a given row, the computer implemented Option A also for all rows above this row. Likewise, if a subject chose Option B in a given row, the computer instantaneously ticked Option B in all lower rows. However, participants could always revise their decision and the auto-completion before moving on. See Figure 12 in Appendix C.1 for a screenshot.

The parameters y and p were drawn uniformly randomly and independently from the sets $y \in \{15, 20, 25\}$ and $p \in \{5, 10, 25, 50, 75, 90, 95\}$. We implemented both gain and loss gambles, where the loss amounts are the mirror images of y. In the case of loss gambles, the lowest safe payment was given by $z = -\$y$ and the highest one by $z = \$0$. In loss choice lists, subjects received a monetary endowment of $\$y$ from which potential losses were deducted. Out of the six choice lists that each subject completed, three concerned loss gambles and three gain gambles. We presented either all loss gambles or all gain gambles first, in randomized order.

Finally, with probability 1/3, a choice list in treatment *Baseline Risk* was presented in a compound lottery format. We will describe, motivate and analyze these data in Section 3.3. For now we focus on the baseline (reduced) lotteries.

3.1.2 Measuring Cognitive Uncertainty

When it comes to measuring cognitive uncertainty about an action, there are two extreme benchmarks. The first is the traditional approach of not measuring it, which amounts to implicitly or explicitly assuming that the decision maker is cognitively certain about the action that he takes. The second benchmark is to elicit the decision-maker's full (probability) distribution around his action. This is tedious in practice. Instead, we resort to measuring a summary statistic that captures the uncertainty implied in the distribution, which is the analog of σ_{CU} in the model (equation (6)). However, many people are not naturally familiar with the concept of a standard deviation. To strike a balance between conceptual clarity and quantitative interpretation on the one hand and participant comprehension on the other hand, we hence elicit an interval measure.

Figure 13 in Appendix C.1 provides a screenshot. Here, a participant was reminded of their valuation (switching interval) for the lottery on the previous price list screen. They were then asked to indicate how certain they are that to them the lottery is worth exactly the same as their previously indicated certainty equivalent. To answer this question, subjects used a slider to calibrate the statement "I am certain that the lottery is worth between a and b to me." If the participant moved the slider to the very right, a and b corresponded to the previously indicated switching interval. For each of the 20 possible ticks that the slider was moved to the left, a decreased and b increased by 25 cents, in real time. In gain lotteries, a was bounded from below by zero and b bounded from above by the lottery's upside. Analogously, for losses, a was bounded from below by the lottery's downside and b from above by zero. The slider was initialized at cognitive uncertainty of zero, but subjects had to click somewhere on the slider in order to be able to proceed.

Four remarks about this measure are in order. First, the measure only captures *internal* uncertainty about what the certainty equivalent is, rather than also external uncertainty that arises due to stochasticity in the environment. Therefore, both traditional and behavioral models that do not feature cognitive noise predict cognitive uncertainty of zero: in canonical models, agents know their true valuation of a lottery, regardless of which "behavioral" decision processes or preferences may generate that valuation.

Second, this measure approximates a subjective confidence interval. Our elicitation procedure did deliberately not specify which particular confidence interval (e.g., 95%) we are interested in. The reason is that (i) we aimed at keeping the elicitation simple and (ii) we are operating precisely under the assumption that subjects do not really know how to translate probabilities of 90% or 95% confidence into an appropriate certainty equivalent. To support this conjecture, Appendix B reports on calibration experiments in which we explicitly elicit 75%, 90%, 95%, 99% and 100% confidence intervals, and compare them with our baseline measure. We find that subjects always state roughly the same cognitive uncertainty ranges, regardless of which confidence interval we elicit.

Third, we deliberately do not financially incentivize our elicitation of cognitive uncertainty. The reason is that we do not know subjects' true valuation in the absence of cognitive noise because we do not know subjects' true preferences. Our approach is therefore related to other recent work in behavioral economics that has highlighted (i) that various non-choice data such as response times or eye-tracking can help better understand choice behavior, and (ii) that unincentivized measurements of behavioral constructs are often highly predictive of behavior (e.g., Falk et al., 2018; Enke et al., 2019), including in the domain of bounded rationality (Stango and Zinman, 2020).

Fourth, our measure of cognitive uncertainty reflects *subjectively perceived* uncertainty about the optimal action, rather than the actual magnitude of cognitive noise. A perhaps intuitively plausible alternative procedure would be to estimate the magnitude of *actual*, latent cognitive noise through across-task inconsistency in behavior. There are four reasons that speak against the usefulness of such a measure in our framework. First, in the model in Section 2 it is actually ambiguous whether higher cognitive noise leads to more inconsistency.[7] The intuition is that higher cognitive noise also increases shrinkage (reduces λ), which can lead to *less* overall variability in behavior. In contrast, the relationship between cognitive uncertainty and both cognitive noise and the degree

[7]From eq. (4), we have $\frac{\partial \sigma^2_{a^o(s)}}{\partial \sigma^2_\varepsilon} = \frac{B^2 \sigma_p^4 (\sigma_p^2 - \sigma_\varepsilon^2)}{(\sigma_p^2 + \sigma_\varepsilon^2)^3}$, which has an ambiguous sign.

of shrinkage is unambiguously positive. A second reason against using across-task inconsistency is that what matters for our logic of Bayesian shrinkage is not necessarily *actual* but *subjectively perceived* noisiness. Third, prior work has shown that people sometimes randomize for reasons that are unrelated to cognitive noise (Agranov and Ortoleva, 2017), or exhibit preferences for behaving consistently (Falk and Zimmermann, 2017). This would confound a measurement based on repeated elicitation. Fourth, we desire a relatively simple measure of cognitive uncertainty at the level of a single price list.

Throughout the paper, we normalize cognitive uncertainty to be in $[0, 1]$, where one corresponds to the widest possible uncertainty interval. Figure 14 in Appendix C.1 shows a histogram of the distribution of cognitive uncertainty. 55% of our data indicate cognitive uncertainty that is strictly larger than the one-dollar switching interval.[8]

3.1.3 Subject Pool

All experiments reported in this paper were conducted on Amazon Mechanical Turk (AMT). AMT is becoming an increasingly used resource in experimental economics (e.g. Imas et al., 2016; DellaVigna and Pope, 2018), including in work on bounded rationality (Martínez-Marquina et al., 2019). Review papers suggest that experimental results on AMT and in the lab closely correspond to each other (Paolacci and Chandler, 2014).

We took four measures to achieve high data quality. First, our financial incentives are unusually large by AMT standards. Average realized earnings in the choice under risk experiments are \$6.10 for a median completion time of 20 minutes. This implies average hourly earnings of \$18, compared to a typical hourly wage of about \$5 on AMT. Second, we screened out inattentive prospective subjects through comprehension questions described below. Third, we pre-registered analyses that remove extreme outliers and speeders. Fourth, subjects only completed six choice lists, which is considerably less than in typical experiments.

3.1.4 Logistics and Pre-Registration

Based on a pre-registration, we recruited $N = 700$ completes for treatment *Baseline Risk*. We restricted our sample to AMT workers that were registered in the United States, but we did not impose additional participation constraints. After reading the instructions, participants completed three comprehension questions. Participants who answered one or more control questions incorrectly were immediately routed out of the experiment and do not count towards the number of completes. In addition, towards the end of the

[8]As a basic validity check, in a small sample of 272 price lists, we implemented payout probabilities of $p = 0\%$ or $p = 100\%$, so that there is no external uncertainty. In these tasks, cognitive uncertainty drops considerably to an average of 0.10 and a median of zero.

experiment, a screen contained a simple attention check. Subjects that answered this attention check incorrectly are excluded from the data analysis and replaced by a new complete, as specified in the pre-registration. In total, 62% of all prospective participants were screened out of the experiment in the comprehension checks. Of those subjects that passed, 2% were screened out in the attention check. These procedures imply that, just like all traditional lab experiments with undergraduates, we are working with a sample that is *positively* selected in terms of cognitive abilities and / or attentiveness. Given the link between cognitive uncertainty, cognitive ability and response times discussed in Section 7, we would probably have identified even more variation in cognitive uncertainty had we not restricted the sample. Screenshots of instructions and control questions can be found in Appendix L.

In terms of timeline, subjects first completed six of the choice under risk tasks discussed above. Then, we elicited their survey expectations about various economic variables, as discussed in Section 5. Finally, participants completed a short demographic questionnaire and an eight-item Raven matrices IQ test.

Participants received a completion fee of \$1.70. In addition, each participant potentially earned a bonus. The experiment comprised three financially incentivized parts: the risky choice lists, the survey expectations questions, and the Raven IQ test. For each subject, one of these parts of the experiment was randomly selected for payment. If choice under risk was selected, a randomly selected decision from a randomly selected choice list was paid out.

The experiments reported in this paper were pre-registered in the AEA RCT registry, see `https://www.socialscienceregistry.org/trials/4493`. The pre-registration includes (i) the sample size in each treatment; (ii) data exclusion criteria such as the aforementioned attention checks or the handling of extreme outliers; and (iii) directional predictions about the relationship between cognitive uncertainty, our outcome measures and experimental manipulations.

3.2 Cognitive Uncertainty and the Probability Weighting Function

Because of the simple structure of our lotteries with only one non-zero payout state, an instructive way to visualize our data is to compute *normalized certainty equivalents* as $NCE = 100 \cdot CE/y$, where the certainty equivalent CE is defined as the midpoint of the switching interval and y is the non-zero payout. An attractive feature of NCE is that it directly corresponds to the implied probability weight if one assumes that utility is linear. For expositional reasons, we change the sign of these normalized certainty equivalents to be weakly negative for loss lotteries.

For the purposes of the baseline analysis, we exclude extreme outliers as defined in the pre-registration: these are observations for which (i) the normalized certainty equivalent is strictly larger than 95% while the objective payout probability is at most 10%, or (ii) the normalized certainty equivalent is strictly less than 5% while the objective payout probability is at least 90%. This procedure affects 3% of all data points. We report robustness checks using all data in Appendix C.2.

Figure 2 plots average normalized certainty equivalents against objective payoff probabilities to visualize the probability weighting function. The figure distinguishes between subjects above and below average cognitive uncertainty within a given probability bucket. Focusing on the upper half of the figure (gain lotteries), first note that we replicate prior findings on the shape of the weighting function. More importantly, we find that subjects with higher cognitive uncertainty exhibit more pronounced probability weighting: high cognitive uncertainty decisions are slightly more risk seeking for small probability gains and more risk averse for high probability gains. Thus, overall, cognitive uncertainty is associated with more pronounced compression.[9]

The heuristic probability weighting function crosses the 45-degree line to the left of $p = 50\%$. This pattern is well-known in the literature and in line with our hypothesis as long as subjects both (i) shrink towards 50:50 because of cognitive noise and (ii) exhibit some version of genuine aversion against risky lotteries.

Next, we turn to the bottom panel of Figure 2, which summarizes the data for loss lotteries. By construction of our figure, the weighting function is now given by the mirror image of the weighting function in the gain domain. Again, we see that the implied probability weights of subjects with higher cognitive uncertainty are more compressed. An attractive feature of visualizing the data as in Figure 2 is that it highlights that the relationship between cognitive uncertainty and risk aversion reverses in predictable ways depending on whether the payouts are positive or negative and whether the payout probability is high or low. For instance, subjects with higher cognitive uncertainty are more risk seeking for small probability gains, but more risk averse for small probability losses. Similarly, high cognitive uncertainty participants are more risk averse for high probability gains, yet more risk seeking for high probability losses. Thus, high cognitive uncertainty subjects exhibit a more pronounced "fourfold pattern of risk attitudes" (Kahneman and Tversky, 1979).

[9]Strictly speaking, our model predicts that probability weighting will entirely disappear in the absence of cognitive noise. In our data, however, we observe some compression also for decisions that are associated with the lowest possible cognitive uncertainty. This pattern is consistent with our model for two reasons. First, because the cognitive uncertainty measure refers to the switching interval in the choice list, the lowest cognitive uncertainty statement in our data does not imply cognitive uncertainty of zero, but rather that cognitive uncertainty does not exceed the \$1 switching interval. Second, because there is plausibly some measurement error in the elicitation of cognitive uncertainty, some decisions will be coded as being associated with the lowest possible level of cognitive uncertainty, while the true level of cognitive uncertainty that governs these decisions is strictly positive. As a result of these two measurement issues, it is unsurprising to see that even "low cognitive uncertainty" decisions reflect some probability weighting.

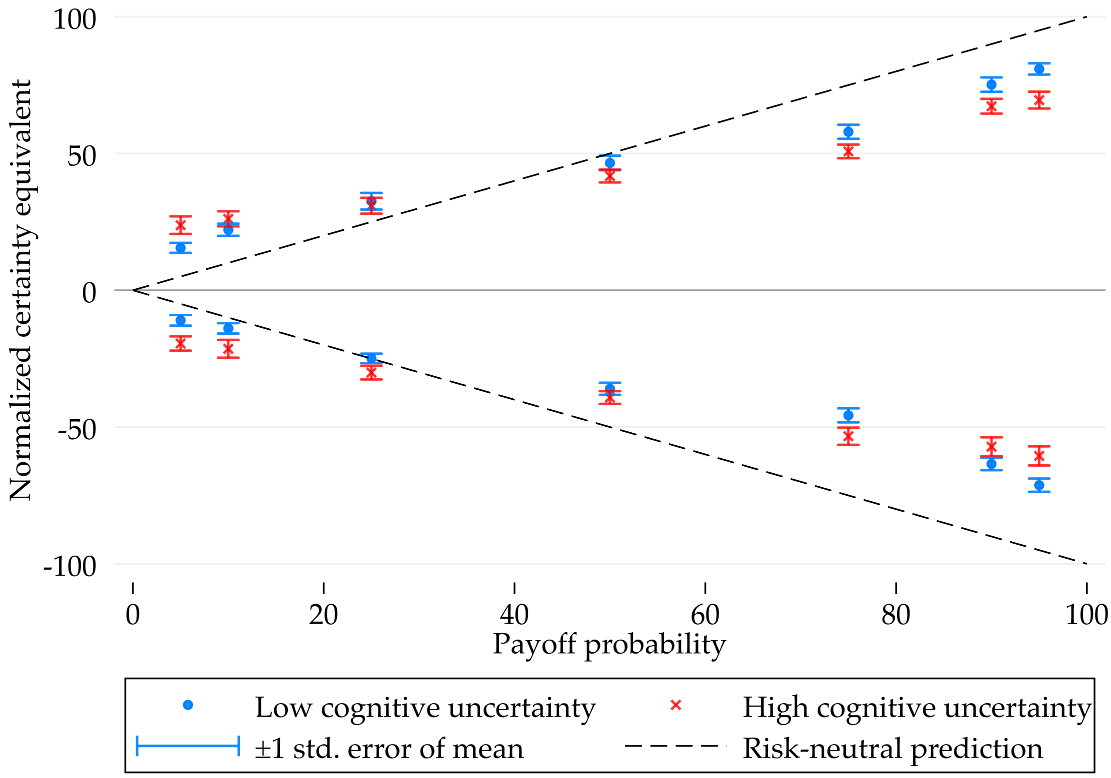

Figure 2: Probability weighting function separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability × gains / losses bucket. The plot shows averages and corresponding standard error bars. The figure is based on 2,525 certainty equivalents of 700 subjects.

Table 1 provides a regression analysis of these patterns, which directly corresponds to estimating the neo-additive weighting function in equation (7). Our object of interest is the extent to which a subject's normalized certainty equivalent is (in)sensitive to variations in the probability of the non-zero payout state. Thus, we regress a participant's absolute normalized certainty equivalent on (i) the probability of receiving the non-zero gain / loss; (ii) cognitive uncertainty; and (iii) an interaction term. In our baseline specification, we do not include subject fixed effects, meaning that we embrace the variation that results from across-subject heterogeneity in cognitive uncertainty.

The results show that higher cognitive uncertainty is associated with lower responsiveness to variations in objective probabilities, in both the gains and the loss domain. In terms of quantitative magnitude, the regression coefficients suggest that with cognitive uncertainty of zero, the slope of the neo-additive weighting function is given by 0.65, yet it is only 0.34 for maximum cognitive uncertainty of one. A different way to gauge quantitative magnitudes is to standardize cognitive uncertainty into a z-score. When doing so, the regression results (not reported) suggest that an one standard deviation increase in cognitive uncertainty decreases the slope of the neo-additive weighting func-

Table 1: Insensitivity to probability and cognitive uncertainty

	Dependent variable: Absolute normalized certainty equivalent						
	Gains		Losses		Pooled		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Probability of payout	0.68***	0.68***	0.59***	0.59***	0.65***	0.65***	0.66***
	(0.02)	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)
Probability of payout × Cognitive uncertainty	-0.41***	-0.41***	-0.20**	-0.19**	-0.31***	-0.31***	-0.26***
	(0.09)	(0.09)	(0.09)	(0.09)	(0.07)	(0.07)	(0.10)
Cognitive uncertainty	11.6**	11.4**	14.8***	14.6***	13.5***	13.9***	9.88*
	(5.19)	(5.27)	(5.26)	(5.25)	(3.84)	(3.87)	(5.85)
Session FE	No	Yes	No	Yes	No	Yes	No
Demographic controls	No	Yes	No	Yes	No	Yes	No
Subject FE	No	No	No	No	No	No	Yes
Observations	1271	1271	1254	1254	2525	2525	2525
R^2	0.54	0.55	0.41	0.42	0.47	0.47	0.66

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's absolute normalized certainty equivalent. The sample includes choices from all baseline gambles with strictly interior payout probabilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

tion by about 0.11. These are arguably large effect sizes that underscore the quantitative relevance of cognitive uncertainty in generating probability weighting.

Column (7) shows that the reduced sensitivity to objective probabilities for high cognitive uncertainty decisions remains statistically highly significant and of a comparable magnitude when we include subject fixed effects in the regression. That is, cognitive uncertainty predicts the magnitude of compression even across decisions made by the same subject. In Section 6, we return to this observation by studying how cognitive uncertainty depends on objective probabilities.

3.3 Manipulations of Cognitive Uncertainty

To exogenously manipulate cognitive uncertainty, we operate with compound lotteries and ambiguous lotteries. To illustrate, consider the case of compound lotteries, where an example lottery is given by: "We randomly draw an integer between 60 and 80, where each number is equally likely to be selected. Call this number n. With probability n%, you receive \$20. With probability 100%-n%, you receive \$0." The corresponding reduced lottery has payout probability $p = 70\%$. These two lotteries are identical under expected utility theory because EU is linear in probabilities. Ambiguous lotteries follow the same format as compound lotteries, except that the distribution from which payoff probabilities are drawn is unknown. An example is: "There is a number *n* that lies between 60 and 80. With probability *n*%, you receive \$20. Otherwise, you receive \$0."

Our hypothesis is that compound and ambiguous lotteries induce higher cognitive uncertainty, which should lead to weighting functions with lower likelihood sensitivity. A causal interpretation of our experiments with respect to cognitive uncertainty requires the assumption that the introduction of compound or ambiguous lotteries affects choices *only* through cognitive uncertainty. While this is a strong assumption, we are not aware of alternative theories that would predict the nuanced pattern of how risk aversion changes as a function of reduced versus compound lotteries, depending on whether the lottery features high or low probabilities and gains or losses.

As noted above, we implemented the compound lotteries as part of treatment *Baseline Risk*, where each lottery had a 1 in 3 chance of being presented in compound form. We collected 1,241 observations on compound lotteries. The ambiguity experiment was added to the pre-registration after the initial set of experiments was implemented. 300 subjects completed these experiments, in which each subject completed both lotteries with known payoff probabilities and ambiguous ones.[10]

Turning to the results, we find that, relative to reduced lotteries, compound and ambiguous lotteries increase stated cognitive uncertainty by 23% and 26%, on average. Figures 15 and 16 in Appendix C.1 show corresponding histograms.

Figure 3 shows the results for the compound manipulation. The analogous figure for ambiguous lotteries is Figure 17 in Appendix C.1. We find that the probability weighting function is substantially more compressed under compound than under reduced lotteries, for both gains and losses. Consistent with many findings in the literature (Halevy, 2007; Gillen et al., 2019), subjects are compound lottery averse for high probability gains (and low probability losses). However, as predicted by our framework, subjects behave as if they are compound risk *loving* for low probability gains and high probability losses. We are not aware of other theories that would predict such a pattern.

Table 2 provides a regression analysis, which again corresponds to estimating the neo-additive weighting function. We find that subjects' certainty equivalents are considerably less responsive to the payout probabilities under compound and ambiguous lotteries than under reduced lotteries, for both gains and losses (in these regressions, the payout probability for ambiguous lotteries is specified as the midpoint of the ambigu-

[10]Appendix J presents an additional ambiguity experiment that we pre-registered and implemented. In these experiments, we do not elicit certainty equivalents for ambiguous lotteries but instead matching probabilities. These experiments also deliver statistically significant evidence for a correlation between cognitive uncertainty and "ambiguity-insensitivity." We relegate these experiments to the appendix both for brevity and, as we discuss in the Appendix, because we believe that they are conceptually more ambiguous than the version that we present in the main text.

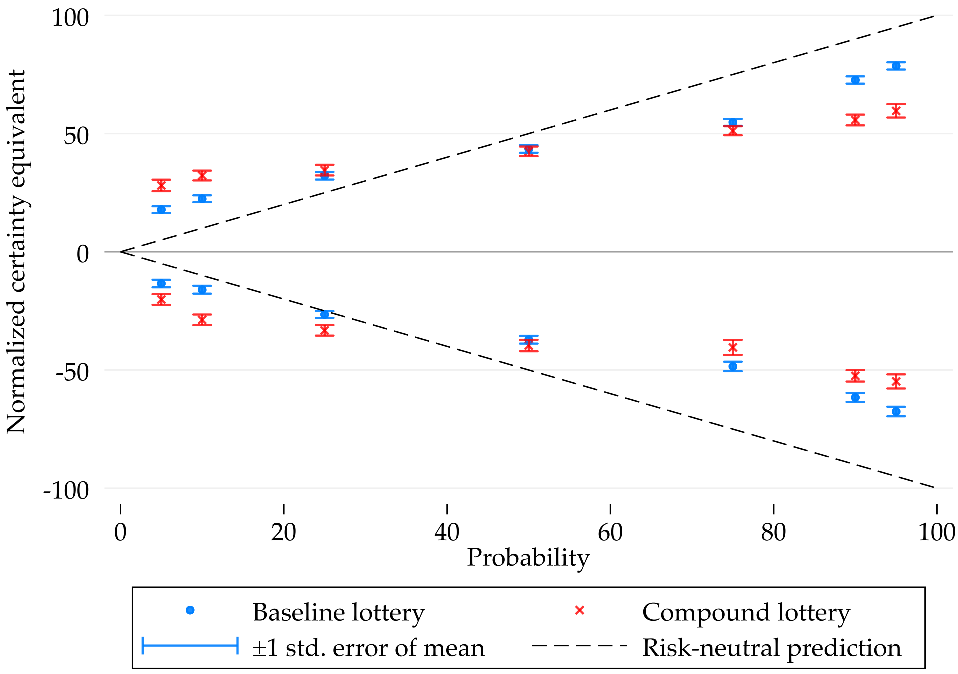

Figure 3: Probability weighting function separately for reduced and compound lotteries. The plot shows averages and corresponding standard error bars. The figure is based on 3,766 certainty equivalents of 700 subjects.

ous interval). Moreover, we again find a within-treatment correlation between responsiveness to payout probabilities and cognitive uncertainty. For example, even when we restrict attention to ambiguous lotteries, the certainty equivalents of participants with higher cognitive uncertainty are significantly less responsive to variation in ambiguous likelihoods than those of subjects with low cognitive uncertainty. This further suggests that the finding of "a-insensitivity" in the ambiguity literature (Trautmann and Van De Kuilen, 2015; Li et al., 2019) reflects cognitive uncertainty. Moreover, we again find that all of our results also hold when we exclusively leverage within-subject variation in cognitive uncertainty, i.e., when subject fixed effects are included in the analyses.

In Appendix K, we report on a further set of experiments in which we exogenously manipulate cognitive uncertainty using cognitive load. In these experiments, we find very similar patterns to those reported above: under cognitive load, stated cognitive uncertainty increases by 15%, and participants' certainty equivalents become less sensitive to variation in objective probabilities, for both gains and losses.

3.4 Manipulation of the Cognitive Default

In a final step of the analysis of choice under risk, we exogenously manipulate the location of the cognitive default. Recall that we operate under the assumption that the

Table 2: Choice under risk: Baseline versus compound / ambiguous lotteries

	Dependent variable: Absolute normalized certainty equivalent					
	Risk vs. compound risk			Risk vs. ambiguity		
	Gains	Losses	Pooled	Gains	Losses	Pooled
	(1)	(2)	(3)	(4)	(5)	(6)
Probability of payout	0.62***	0.56***	0.65***	0.74***	0.54***	0.69***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.04)	(0.03)
Probability of payout × 1 if compound / ambiguous lottery	-0.30***	-0.25***	-0.29***	-0.20***	-0.16***	-0.16***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.02)
Probability of payout × Cognitive uncertainty			-0.24***			-0.39***
			(0.07)			(0.09)
1 if compound / ambiguous lottery	12.3***	12.3***	13.6***	6.91***	8.82***	6.72***
	(1.89)	(1.84)	(1.60)	(1.14)	(2.31)	(1.35)
Cognitive uncertainty			10.5**			16.2***
			(4.12)			(5.66)
Subject FE	No	No	Yes	No	No	Yes
Observations	1918	1848	3766	889	880	1769
R^2	0.44	0.35	0.56	0.58	0.34	0.71

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's absolute normalized certainty equivalent. In columns (1)–(3), the sample includes choices from the baseline and compound lotteries, where for comparability the set of baseline lotteries is restricted to lotteries with payout probabilities of 10%, 25%, 50%, 75%, and 90%, see Figure 3. In columns (4)–(6), the sample includes choices from the baseline and ambiguous lotteries. For ambiguous lotteries, we define the payout "probability" as the midpoint of the interval of possible payout probabilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

default is influenced by an ignorance prior. With two states of the world, the ignorance prior is 50:50. To vary the default, we implement a partition manipulation (Starmer and Sugden, 1993; Fox and Clemen, 2005) and increase the number of states to ten. This means that the ignorance prior for each state is now given by 10%. We further designed this treatment variation with the objective of holding cognitive uncertainty fixed (which we verify below). Following the logic of Prediction 3 in Section 2, we predict that the elevation of the probability weighting function decreases as the number of states increases.

To experimentally implement this manipulation, we replicate treatment *Baseline Risk*, but now frame probabilities in terms of number of colored balls in an urn. For example, we describe a lottery as:

> Out of 100 balls, 80 are red. If a red ball gets drawn: get $20.
> 20 balls are blue. If a blue ball gets drawn: get $0.

In addition to this treatment, labeled *High Default Risk*, we also implement treatment *Low Default Risk*. Here, we implement the same lotteries as in *High Default Risk*, yet we split the zero-payout state into nine payoff-equivalent states with different probability colors. For example, the lottery above would be described as:

> Out of 100 balls, 80 are red. If a red ball gets drawn: get \$20.
> 2 balls are blue. If a blue ball gets drawn: get \$0.
> 2 balls are black. If a black ball gets drawn: get \$0.
> 2 balls are white. If a white ball gets drawn: get \$0.
> ...
> 4 balls are yellow. If a yellow ball gets drawn: get \$0.

It is important to note that these lotteries are *identical* in terms of objective payout profiles. Still, based on our model, we predict (i) that this manipulation shifts the probability weighting function towards zero and (ii) that high cognitive uncertainty decisions are more responsive to variation in the cognitive default. In total, 300 subjects participated in these two treatments, which we implemented in a between-subjects design with random assignment to treatments within sessions.

Turning to the results, we find that cognitive uncertainty does not vary across the two treatments ($p = 0.898$), see the histograms in Figure 18 in Appendix C.1. This lends credence to our implicit assumption that our experimental manipulation only affects the cognitive default but not cognitive uncertainty.

Figure 4 shows average normalized certainty equivalents, separately for treatments *High Default Risk* and *Low Default Risk*. We find that, in the gain domain, the probability weighting function is significantly shifted downwards towards zero with 10 states (a low default), as hypothesized. In the loss domain, our framework would predict that the weighting function is shifted upwards towards zero. We only find mixed evidence for this prediction: the weighting function appears to move up for low and intermediate probabilities but not for high probabilities.

Table 3 provides a regression analysis that confirms the visual patterns. Because the treatment variations were conducted across subjects, we cannot include subject fixed effects in the regressions. Columns (1)–(2) analyze gain lotteries. Here, normalized certainty equivalents (observed risk tolerance) are 10 percentage points lower in the *Low Default Risk* condition. Moreover, as we can see in column (2), we find a statistically significant interaction effect between cognitive uncertainty and the treatment manipulation. This suggests that, as posited by our model, subjects with high cognitive noise respond more to variation in the cognitive default. In the case of losses, the regression coefficient of the low default condition is negative – as predicted by our framework – but not statistically significant ($p = 0.15$). The same is true for the hypothesized interaction

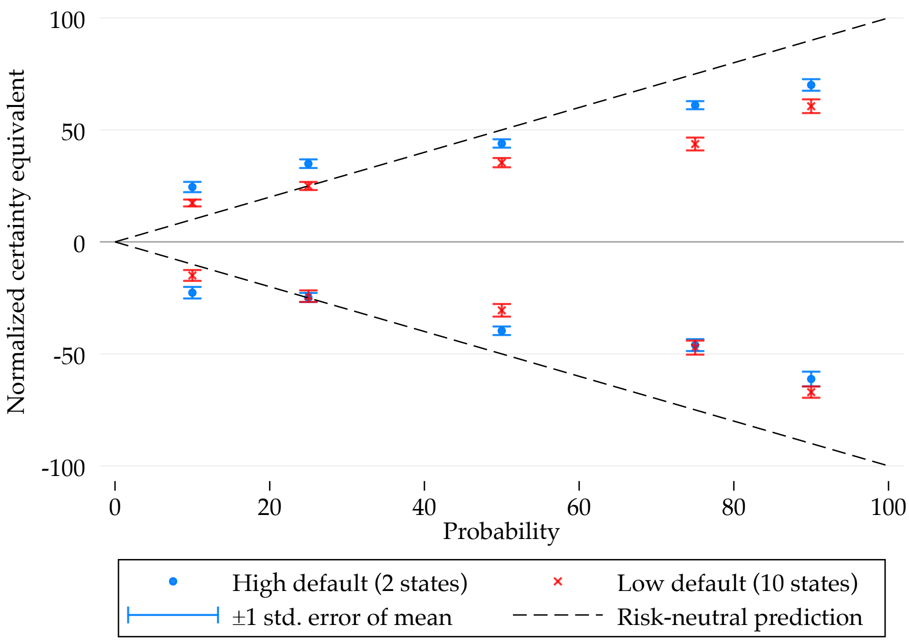

Figure 4: Probability weighting function separately for treatments *High Default Risk* and *Low Default Risk*. The plot shows averages and corresponding standard error bars. The figure is based on 1,757 certainty equivalents of 300 subjects.

coefficient. A potential (post-hoc) explanation for this null result is that, in all treatments, the choice data in the loss domain appear to be considerably noisier than in the gain domain. This can be inferrred from the difference in R^2 between columns (1) and (3) in Table 3 and similar patterns in all other tables above. Either way, both the treatment effect of the default manipulation and its interaction with cognitive uncertainty are statistically significant in the pooled gains and losses sample, see columns (5)–(7).

4 Belief Updating

4.1 Experimental Design

Our experimental design strategy for belief updating closely mirrors the one for choice under risk: we (i) supplement an established experimental design from the literature with a measurement of cognitive uncertainty; (ii) document a correlation between cognitive uncertainty and the magnitude of compression of probabilities; (iii) exogenously manipulate cognitive uncertainty using a compound manipulation; and (iv) vary the location of the cognitive default using a partition manipulation.

Table 3: Choice under risk: Treatments *Low Default Risk* and *High Default Risk*

	Dependent variable: Absolute normalized certainty equivalent						
	Gains		Losses		Pooled		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
0 if *High Default*, 1 if *Low Default*	-10.4***	-6.70***	-1.03	-2.33	-6.46***	-4.37**	-3.67*
	(1.84)	(2.27)	(2.45)	(2.88)	(1.50)	(1.86)	(1.87)
Probability of payout	0.52***	0.52***		0.52***	0.53***	0.53***	0.53***
	(0.03)	(0.03)		(0.03)	(0.02)	(0.02)	(0.02)
1 if *Low Default* × Cognitive uncertainty		-18.6**		-0.64		-10.2*	-11.7**
		(7.55)		(9.14)		(5.87)	(5.95)
Cognitive uncertainty		4.63		10.6*		7.74*	8.18*
		(6.22)		(5.84)		(4.17)	(4.26)
Session FE	No	No	No	No	No	No	Yes
Demographic controls	No	No	No	No	No	No	Yes
Observations	881	881	876	876	1757	1757	1757
R^2	0.39	0.40	0.00	0.30	0.33	0.33	0.34

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's absolute normalized certainty equivalent. The sample includes choices from treatments *Low Default Risk* and *High Default Risk*. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.1.1 Measuring Belief Updating

In designing a structured belief updating task, we follow the recent review and meta-study by Benjamin (2019) by implementing the workhorse paradigm of so-called "balls-and-urns" or "bookbags-and-pokerchips" experiments. In treatment *Baseline Beliefs*, there are two bags, A and B. Both bags contain 100 balls, some of which are red and some of which are blue. The computer randomly selects one of the bags according to a pre-specified base rate. Subjects do not observe which bag was selected. Instead, the computer selects one or more of the balls from the selected bag at random (with replacement) and shows them to the subject. The subject is then asked to state a probabilistic guess that either bag was selected. We visualized this procedure for subjects using the image at the top right of Figure 20 in Appendix D.1.

The three key parameters of this belief updating problem are: (i) the base rate $r \in \{10, 30, 50, 70, 90\}$ (in percent), which we operationalized as the number of cards out of 100 that had "bag A" or "bag B" written on them; (ii) the signal diagnosticity $q \in \{70, 90\}$, which is given by the number of red balls in bag A and the number of blue balls in bag B (we only implemented symmetric signal structures such that $P(red|A) = P(blue|B)$); and (iii) the number of randomly drawn balls $N \in \{1, 3\}$. These parameters were randomized across trials but always known to participants.

Each subject completed six belief updating tasks. In each task, they were asked to state a probabilistic belief (0-100) that bag A got selected. The computer automatically and instantaneously showed the corresponding subjective probability that bag B got selected. See Figure 19 in Appendix D.1 for a screenshot.

Financial incentives were implemented through the binarized scoring rule (Hossain and Okui, 2013). Here, subjects had a chance of winning a prize of \$10. The probability of receiving the prize was given by $\pi = \max\left\{0, \frac{100-0.04\cdot(b-t)^2}{100}\right\}$, where b is the guess (in %) and t indicates the true state (0 or 100).[11]

With probability 5 in 6, a belief updating task was implemented using the design discussed above, and with probability 1 in 6 in a compound design. We return to the compound data in Section 4.3 and focus on the baseline problems for now.

4.1.2 Measuring Cognitive Uncertainty

Our main measure of cognitive uncertainty in belief updating is very similar to the one for choice under risk, both conceptually and implementation-wise. The instructions introduced the concept of an "optimal guess." This guess, we explained to subjects, uses the laws of probability to compute a statistically correct statement of the probability that either bag was drawn, based on Bayes' rule. We highlighted that this optimal guess does not rely on information that the subject does not have.

After subjects had indicated their probabilistic belief that either bag was drawn, the next decision screen elicited cognitive uncertainty. Here, we asked subjects how certain they are that their own guess equals the optimal guess for this task. Operationally, similarly to the case of choice under risk, subjects navigated a slider to calibrate the statement "I am certain that the optimal guess is between a and b.", where a and b collapsed to the subject's own previously indicated guess in case the slider was moved to the very right. For each of the 30 possible ticks that the slider was moved to the left, a decreased and b increased by one percentage point. a was bounded from below by zero and b bounded from above by 100. Again, the slider was initialized at cognitive uncertainty of zero and we forced subjects to click somewhere on the slider to be able to proceed. Figure 20 in Appendix D.1 shows a screenshot of the elicitation screen. For ease of interpretation, we again normalize this measure to be between zero and one. As in choice under risk, this measure only captures *internal* uncertainty about what the rational solution to the decision problem is, rather than stochasticity in the environment.

Just like our measure of cognitive uncertainty in choice under risk, this one is not financially incentivized. However, in the case of belief updating, it is possible to devise an incentivized measure because here an objectively optimal response (the Bayesian

[11]Recent evidence suggests that biases in probability judgment are largely invariant to the stake size (Enke et al., 2020).

posterior) exists. Thus, we additionally elicited a second measure of cognitive uncertainty from each participant: their willingness-to-pay (WTP) for replacing their own guess with the optimal (Bayesian) guess.[12] To this effect, before subjects stated their own guess, they received an endowment of \$3 for each task and then indicated how much of this budget they would at most be willing to pay to replace their guess.[13] Subjects' WTP was elicited using a direct Becker-deGroot-Marschak elicitation mechanism. That is, we randomly drew a price $p \sim U[0,3]$ and the guess was replaced iff $p \leq$ WTP. See Figure 21 in Appendix D.1 for a screenshot.

To maximize statistical power, subjects' WTP and the resulting replacement of their own decision was only implemented in randomly selected 10% of all tasks. To avoid concerns about hedging, this uncertainty was resolved before subjects stated their own posterior guess. The timeline of each task was hence as follows: (i) observe game parameters; (ii) indicate WTP; (iii) find out whether own guess or Bayesian guess potentially counts for payment; (iv) state own posterior guess; and (v) indicate cognitive uncertainty range. The analysis below excludes those tasks in which a subject's guess got replaced by the optimal guess (3% of all data), though we have verified that our analyses yield virtually identical results if these (non-incentivized) guesses are included.

Figures 22 and 23 in Appendix D.1 show histograms of the cognitive uncertainty measure as well as subjects' WTP. Both measures exhibit considerable variation. 85% of our data indicate strictly positive cognitive uncertainty. The average WTP is \$0.85 with a median of \$0.50 and a standard deviation of 0.93.[14]

The two measures exhibit a correlation of $\rho = 0.21$. While not incentivized, we view the cognitive uncertainty measure as our primary measure because (i) by its nature, and as exemplified by this paper, it is easily portable across different experimental contexts and decision situations and (ii) it is more fine-grained and exhibits more variation (26% of all WTPs are zero, perhaps due to loss aversion vis-a-vis giving up safe money). Still, below we verify that all of our results are robust to using the WTP measure.

[12]See Oprea (2020) for a related recent approach in the context of algorithmic rules.

[13]To see the intuition for why even risk-neutral subjects who have strictly positive cognitive uncertainty but are otherwise rational would have a non-zero WTP for the Bayesian posterior, consider the following contrieved example. Suppose the signal diagnosticity is one, i.e., signals are perfectly informative. The subject states posterior $p = 0.5$, but actually believes that $p = 0$, $p = 0.5$ and $p = 1$ are all equally likely to be the Bayesian posterior. Then, purchasing the Bayesian guess amounts to perfectly learning the true state, while guessing 0.5 evidently has lower expected payoffs.

[14]As a basic validity check, in a small sample of 161 updating tasks, we implemented a signal diagnosticity of $d = 100$, so that the selected bag is deterministically revealed. In these tasks, the distribution of both the cognitive uncertainty range and subjects' WTP has a median of zero, with means of 0.06 and 0.26.

4.1.3 Logistics and Pre-Registration

Based on a pre-registration, we recruited $N = 700$ completes for treatment *Baseline Beliefs*. Participants who answered one or more of the four comprehension questions incorrectly were immediately routed out of the experiment. Similarly, subjects are excluded from the analysis if they failed an attention check, as specified in the pre-registration. In total, 49% of all prospective participants were screened out in the comprehension checks. Of those subjects that passed, 6% were screened out based on the attention check.

In terms of timeline, subjects first completed the belief updating tasks discussed above. Second, we elicited their survey expectations about economic variables, discussed in Section 5. Finally, participants completed a short demographic questionnaire and an eight-item Raven matrices IQ test. One of the three parts of the experiments (belief updating, survey expectations, or Raven test) was randomly selected for payment.

Average earnings are \$4.80 with a median completion time of 23 minutes. The experiments were pre-registered under the same AEA RCT trial as discussed above. Screenshots of the instructions and control questions can be found in Appendix L.

4.2 Cognitive Uncertainty and Belief Updating

As in the analysis of choice under risk and as specified in the pre-registration, we begin by excluding extreme outliers. These are defined for subjective posteriors p_s and Bayesian posteriors p_b such that $p_s < 25 \wedge p_b > 75$ or $p_s > 75 \wedge p_b < 25$. This is the case for 5% of all data. We report robustness checks using the full sample below.

Figure 1 in the Introduction depicts the "belief weighting function" that we estimate in our data: the inverse S-shaped relationship between average stated and Bayesian posteriors that is also documented in Ambuehl and Li (2018). Figure 5 replicates this figure separately for subjects above or below average cognitive uncertainty as defined by our unincentivized cognitive uncertainty range. We see that, over the entire support of Bayesian posteriors, stated posteriors are more compressed towards 50:50 for subjects with higher cognitive uncertainty.[15] Thus, cognitive uncertainty endogenizes the well-known phenomenon of "extremeness aversion" in belief updating. Figure 27 in Appendix D.2 replicates this figure based on the financially incentivized WTP measure, with very similar results.

Columns (1)–(3) of Table 4 provide an econometric analysis, which again corresponds to the neo-additive weighting function. Here, we regress a subject's stated pos-

[15] A commonly observed error in belief updating experiments is that subjects report the signal diagnosticity. In our data, there are 4,289 belief statements for tasks in which the Bayesian posterior did not equal the signal diagnosticity. Out of those, 374 (8.7%) directly reflect the signal diagnosticity. All of our results are robust to excluding these observations.

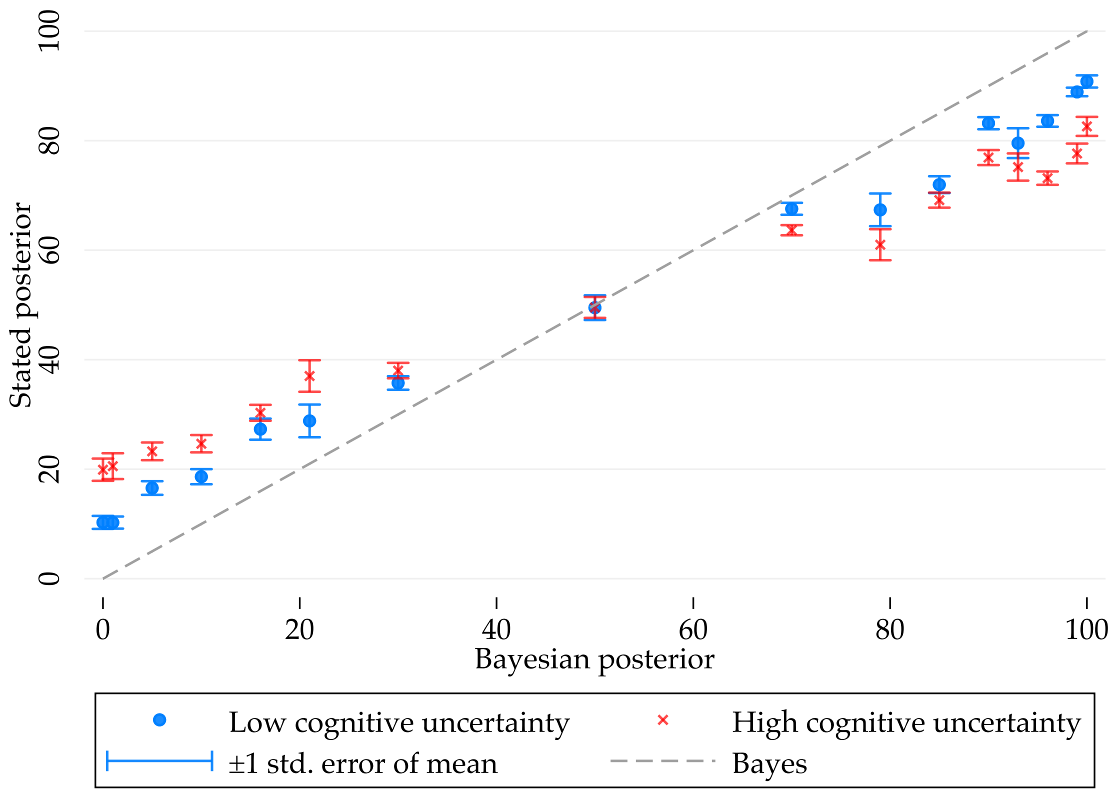

Figure 5: Relationship between average stated and Bayesian posteriors, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each Bayesian posterior. Bayesian posteriors are rounded to the nearest integer. We only show buckets with at least 15 observations. The figure is based on 3,187 beliefs of 700 subjects.

terior on (i) the Bayesian posterior; (ii) cognitive uncertainty; and (iii) their interaction term. We find that with cognitive uncertainty of zero, the slope of the neo-additive weighting function is given by 0.83 but it is only 0.41 with cognitive uncertainty of one. Put differently, our data suggest that a one standard deviation increase in cognitive uncertainty reduces the slope of the linear function by 0.10.[16]

Insensitivity to base rate and likelihood ratio (conservatism). A canonical finding in the belief updating literature is that subjects exhibit underreaction (insensitivity) to both the base rate and the likelihood ratio, which is also referred to as conservatism (Benjamin, 2019). Note that our theory of cognitive uncertainty and belief updating mechanically generates insensitivity to the base rate and the likelihood ratio because – in our model – subjects only respond to variation in the Bayesian posterior (which is determined by the base rate and the likelihood ratio) with weight $\lambda < 1$. Intuitively, if a subject always states a belief of 50:50, the sensitivity of beliefs to the base rate and like-

[16]Table 11 in Appendix D.2 replicates Table 4 using the WTP instead of the cognitive uncertainty measure, with very similar results.

Table 4: Belief updating: Regression analyses

	Dependent variable: Posterior belief								
Sample:	Baseline			Compound			Default		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Bayesian posterior	0.80*** (0.01)	0.80*** (0.01)	0.80*** (0.02)	0.72*** (0.02)	0.80*** (0.02)	0.80*** (0.03)	0.64*** (0.01)	0.64*** (0.01)	0.64*** (0.01)
Bayesian posterior × Cognitive uncertainty	-0.39*** (0.04)	-0.39*** (0.04)	-0.39*** (0.05)		-0.28*** (0.05)	-0.26*** (0.07)			
Cognitive uncertainty	16.6*** (2.32)	16.4*** (2.32)	15.9*** (3.09)		10.4*** (3.02)	7.82* (4.62)			
Bayesian posterior × 1 if compound problem				-0.51*** (0.03)	-0.47*** (0.03)	-0.49*** (0.04)			
1 if compound problem				26.4*** (1.75)	25.4*** (1.77)	26.8*** (2.51)			
0 if *Baseline*, 1 if *Low Default*							-3.75*** (0.71)	-1.94* (1.07)	-2.01* (1.14)
1 if *Low Default* × Cognitive uncertainty								-5.78*** (2.15)	-5.99*** (2.17)
Session FE	No	Yes	No	No	Yes	No	No	No	Yes
Demographic controls	No	Yes	No	No	Yes	No	No	No	Yes
Subject FE	No	No	Yes	No	No	Yes	No	No	No
Observations	3187	3187	3187	1947	1947	1947	5372	5372	5372
R^2	0.73	0.74	0.80	0.60	0.61	0.77	0.63	0.63	0.63

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. In columns (1)–(3), the sample includes the baseline tasks. In columns (4)–(6), the sample includes the baseline and compound tasks (where the sample of baseline tasks is restricted to the same probabilities as in the compound tasks). In columns (7)–(9), the sample includes the low and high default tasks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

lihood ratio is zero. To show this empirically, we resort to so-called Grether regressions (Grether, 1980). This specification is derived by expressing Bayes' rule in logarithmic form, which implies a linear relationship between the posterior odds, the prior odds, and the likelihood ratio. Table 12 in Appendix D.3 presents such Grether regressions. We find that insensitivity to the base rate and to the likelihood ratio are indeed significantly more pronounced for subjects with higher cognitive uncertainty. These patterns suggest that (at least a part of) what this literature has identified as base rate neglect or conservatism are in fact not independent psychological phenomena but instead generated by cognitive noise.

Sample size effects. At the most basic level, our account of cognitive uncertainty provides a mapping from Bayesian posteriors to stated posteriors. As discussed above, this endogenizes "extreme belief aversion" (Benjamin, 2019), base rate insensitivity and conservatism. However, as is well known in the literature, experimental data also reveal systematic variation in stated beliefs *conditional* on Bayesian posteriors. For instance, for a given base rate, the draw of one blue ball gives rise to the same Bayesian posterior as the draw of two blue balls and one red ball, yet experimental participants consistently update more strongly after observing one blue ball (Benjamin, 2019). The commonly advanced intuition is that subjects update based on sample *proportions* (which are more extreme for smaller samples), while Bayesian updating prescribes updating based on sample *differences*. Our account of cognitive uncertainty also provides an explanation for this pattern. The straightforward reason is that, in our data, stated cognitive uncertainty significantly increases in the sample size, holding the sample difference fixed (see Table 13 in Appendix D.4). That is, subjects appear to find it easier to form beliefs based on one blue ball than based on two blue balls and one red ball. As a result of this systematic variation in cognitive noise, our model correctly predicts that subjects respond more when the sample size is smaller (and therefore the sample proportion more extreme).

4.3 Manipulation of Cognitive Uncertainty

To manipulate cognitive uncertainty, we again resort to turning baseline problems into compound problems. Consider belief updating problems in which the base rate is given by 50:50 and the signal diagnosticity by $d \equiv P(A|red) = P(B|blue)$. In the compound version of these problems, the base rate is known and fixed at 50:50, but the diagnosticity is the outcome of a random draw, $k \sim U\{d-10, d-9, \ldots, d+10\}$. It is straightforward to verify that these two problems give rise to the same Bayesian posterior.

As in choice under risk, we hypothesize that subjects exhibit higher cognitive uncertainty in compound than in reduced updating problems. By the logic of our framework, we expect that participants' beliefs in compound problems will be more compressed towards 50:50.[17]

We implemented these compound belief updating problems as part of treatment *Baseline Beliefs*, where each belief updating problem had a 1 in 6 chance of being presented in a compound form. We collected 592 observations on compound problems.

We find that, relative to reduced updating problems, compound signal diagnosticities increase stated cognitive uncertainty by 33% and subjects' WTP for the Bayesian guess by 43%, on average. Figures 24 and 25 in Appendix D.1 show corresponding histograms.

Figure 6 summarizes the results. Here, we plot average stated posteriors as a function of Bayesian posteriors, separately for baseline and compound problems. Because in compound problems the base rate is always 50:50, the figure only includes data from problems with a 50:50 base rate for the baseline tasks, too. We see that subjects' posteriors are substantially more compressed towards 50:50 in compound updating problems. Columns (4)–(6) of Table 4 provide a regression analysis. The regression coefficients sug-

[17]In contemporaneous work, Liang (2019) identifies underreaction under compound relative to reduced updating problems. This is in line with our work, but he does not measure cognitive uncertainty.

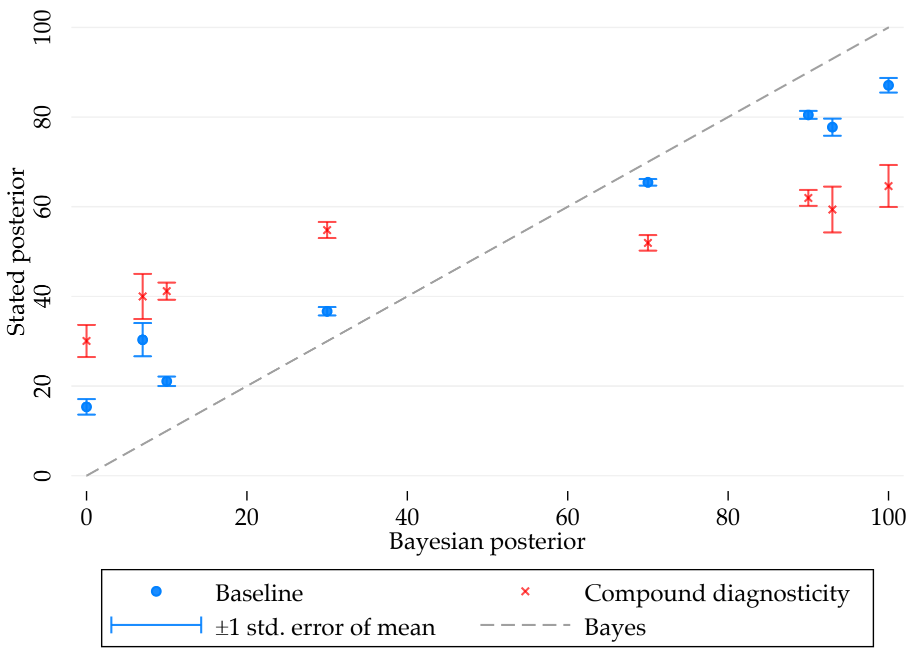

Figure 6: Stated average posteriors as a function of Bayesian posteriors, separately for reduced and compound belief updating problems. The plot shows averages and corresponding standard error bars. To allow for a valid comparison between baseline and compound updating problems, the sample is restricted to updating tasks in which the base rate is 50:50. Bayesian posteriors are rounded to the nearest integer. We only show buckets with at least 15 observations. The figure is based on 1,947 beliefs of 691 subjects.

gest that the sensitivity of stated posteriors to the Bayesian posterior is 0.81 in baseline updating problem, yet only 0.35 in compound updating problems.

4.4 Manipulation of the Cognitive Default

To manipulate the location of the cognitive default, we again employ a partition manipulation and increase the number of states to ten. Under our maintained assumption that the default is influenced by an ignorance prior, our framework predicts that the entire distribution of posterior beliefs shifts towards zero.

Recall that in treatment *Baseline Beliefs*, an example updating problem is that the base rates for bags A and B are 70% and 30%, and the signal diagnosticity (number of red balls in bag A and number of blue balls in bag B) 70%. Now, in treatment *Low Default Beliefs*, we split the probability mass for bag B up into nine different bags. That is, there are now ten bags, labeled A through J. In the specific example above, the base rate for A would again be 70%, the one for B through I 3% each and the one for J 6%. Bag A would contain 70 red and 30 balls, and all bags B through J 30 red and 70 blue balls. That is, these bags have identical ball compositions.

Note that, regardless of what the actual draws of balls are, the Bayesian posterior for bag A having been selected is identical in the baseline version and the version with

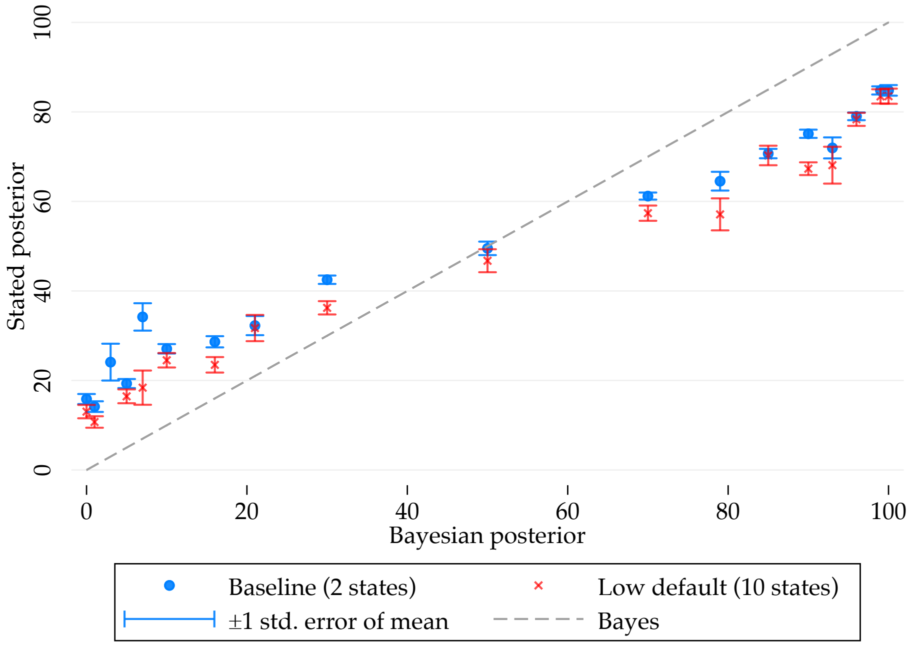

Figure 7: Stated average posteriors as a function of Bayesian posteriors, separately for treatments *Baseline Beliefs* and *Low Default Beliefs*. Bayesian posteriors are rounded to the nearest integer. We only show buckets with at least 15 observations. The figure is based on 5,372 beliefs of 1,000 subjects.

10 bags. The reason is that under the state space {A; not A} the base rates and signal diagnosticities are identical. Thus, in treatment *Low Default Beliefs*, we asked participants to indicate their belief that bag A got selected, and the computer automatically showed the corresponding composite probability for one of the other bags having been selected.

300 subjects participated in treatment *Low Default Beliefs*, which was randomized within the same experimental sessions as treatment *Baseline Beliefs*. All procedures other than the ones described above were identical to the ones in *Baseline Beliefs*. Based on our model, we predict (i) that the entire distribution of posterior beliefs shifts towards zero and (ii) that high cognitive uncertainty decisions are more sensitive to the treatment variation of the cognitive default.

We find that stated cognitive uncertainty is almost identical across conditions *Baseline Beliefs* and *Low Default Beliefs* ($p = 0.85$). This corroborates our implicit assumption that the experimental manipulation of increasing the number of bags only manipulates the cognitive default but not cognitive uncertainty.

Figure 7 shows average stated posteriors as a function of Bayesian posteriors, separately for treatments *Baseline Beliefs* and *Low Default Beliefs*. As predicted, the entire distribution of subjects' beliefs is shifted towards zero. Columns (5)–(7) and (12)–(14) of Table 4 provide regression analyses that confirm the visual patterns by showing that the treatment dummy for the *Low Default* treatment is quantitatively large and statistically significant.

Moreover, the regressions also confirm our second hypothesis, which is that the treatment effect of the cognitive default variation is more pronounced among high cognitive uncertainty subjects (see columns (7)–(9)). Overall, these results highlight that cognitive uncertainty does not produce a general form of insensitivity. Instead, as predicted by our model, it produces insensitivity to objective probabilities but excessive sensitivity to variation in the normatively irrelevant cognitive default.

5 Survey Expectations

5.1 Experimental Design

Unlike the choice and belief updating contexts discussed so far, survey forecasts of economic variables rely on information acquired outside of the experimental context. However, our concept of cognitive uncertainty nevertheless applies here: people oftentimes do not know the answer to a probabilistic question, which may induce them to shrink their reported beliefs to 50:50 (Lebreton et al., 2015). To illustrate the link between cognitive uncertainty and survey expectations, we elicit beliefs about three variables that have attracted attention in the literature: the structure of the national income distribution, inflation rates, and the development of the stock market.

To financially incentivize responses without going through the logistical hassle of waiting for future variables to have realized, we elicited beliefs about contemporaneous or past variables. Each participant was asked three questions that elicited their beliefs about some specific aspect of the income distribution, stock returns, and the inflation rate. The question about the inflation rate reads as:

> *[Explanation of inflation rates.] We randomly picked a year X between 1980 and 2018. Imagine that, at the beginning of year X, the set of products that is used to compute the inflation rate cost \$100. What do you think is the probability that, at the end of that same year, the same set of products cost less than \$y? (In other words, what do you think is the probability that the inflation rate in year X was lower than z%?)*

Beliefs about the income distribution in the United States and the performance of the S&P 500 stock market index were elicited in similar ways; see Appendix E.1.

The order of topics was randomized across participants. Across participants, y (and hence z) varies randomly such that the true probability ranges fom 0% to 100%. Subjects' beliefs were financially incentivized using the same binarized scoring rule as discussed in Section 4, except that the prize a subject could win was \$2. One of the three questions was randomly selected for payment.

To measure cognitive uncertainty, we make use of the same elicitation tool as before. That is, subjects were asked how certain they are that their probabilistic guess is correct. Subjects used a slider to calibrate the statement: "I am certain that the actual probability that [...] is between a and b.", where a and b collapsed to the subject's own previously indicated guess if the slider was moved to the very right. For each of the 30 possible ticks that the slider was moved to the left, a decreased and b increased by one probability point. Figure 28 in Appendix E.2 shows a screenshot of the elicitation screen.

The elicitation of survey expectations took place with the same set of subjects that completed the choice under risk and belief updating tasks discussed in Sections 3 and 4. Thus, the total sample size is $N = 2,000$. Figures 29–31 in Appendix E.2 show histograms of cognitive uncertainty for each question type.

In addition to these "backward-looking" beliefs, in separate pre-registered robustness experiments with $N = 400$ participants, we also elicit expectations about *future* realizations of inflation rates, stock market returns and the income distribution. These questions are conceptually closer to "expectations" in that they ask about the future, but they are not financially incentivized. The results in these robustness experiments are almost identical to the ones that are reported here; we summarize them in Appendix F.

5.2 Results

As in Section 4 and according to our pre-registration, we begin by excluding extreme outliers, defined as $p_s < 25 \wedge p_b > 75$ or $p_s > 75 \wedge p_b < 25$, where p_s is the subjective probability and p_b the objectively correct one. This results in the exclusion of 5% of all data. Figure 8 shows average beliefs as a function of objective probabilities, separately for subjects with above and below average cognitive uncertainty.[18] To conserve space, we only show the results for inflation and stock market expectations; the results for income distribution beliefs are similar and shown in Figure 32 in Appendix E.2. We see that stated beliefs are compressed towards 50%, and that this pattern is substantially more pronounced for subjects who indicate higher cognitive uncertainty. Table 16 in Appendix E.3 provides a corresponding econometric analysis that confirms the statistical significance of this pattern.

6 On Inverse S-Shapes

Our analyses thus far rely on linear regressions that are directly motivated by the neo-additive weighting function that we endogenized in Section 2. While these linear representations are popular due to their simplicity, they have the drawback that they do not

[18]"Objective probabilities" are backward-looking and correspond to the period 1980–2018.

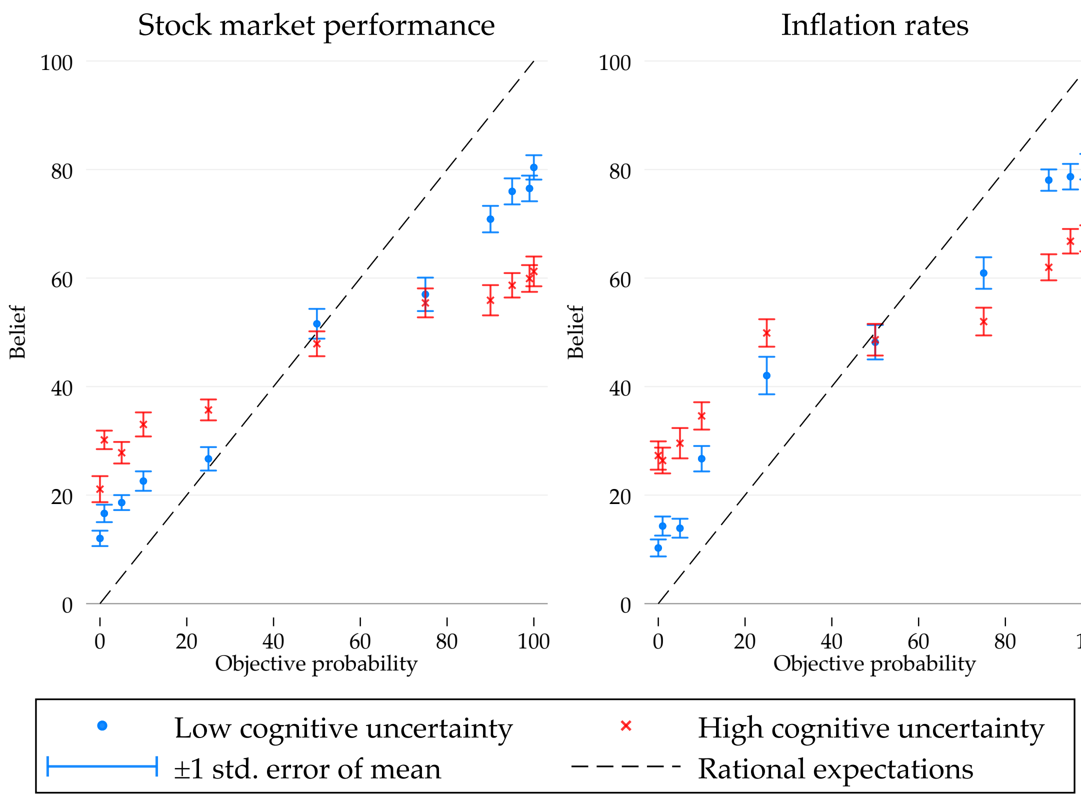

Figure 8: Survey beliefs as a function of objective probabilities, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability bucket. In the left panel, the question asks for the probability that in a randomly selected year the S&P500 increased by less than x% ($N = 1,887$). In the right panel, the question asks for the probability that in a randomly selected year the inflation rate was less than x% ($N = 1,842$).

fully capture the canonical inverse S-shaped response patterns summarized in Figure 1. Inverse S-shaped functions are steeper close to the boundaries of zero and one than in the intermediate range of about $p \in [0.25, 0.75]$. Linear fits of the data generate predicted values that lie (i) above the observed data for small probabilities close to zero and (ii) below the observed data for high probabilities close to one.

The upshot of this discussion is that – in relative terms – there appears to be *more compression towards 50:50 for intermediate than for extreme probabilities*. Indeed, in all of our applications in Sections 3–5, joint Wald tests strongly reject the assumption that objective probabilities and implied subjective probabilities are linearly related. This section aims at documenting that cognitive uncertainty explains not only the general pattern of shrinkage towards 50:50 but also the non-linearities in this process, i.e., the different degrees of compression in different probability ranges.

6.1 Variable Cognitive Noise Across the Probability Range

The formal framework in Section 2 implicitly assumes that the variance of cognitive noise σ_ε^2 is constant in p. While this assumption is analytically convenient, it is strongly and systematically rejected in the data. The three left-hand panels of Figure 9 show the empirical relationship between measured cognitive uncertainty and objective probabilities. In the top left panel, the x-axis shows the objective payout probability of a gamble. In the center left panel, the x-axis denotes the Bayesian posterior in belief updating tasks and in the lower left panel the objectively correct probability in survey estimation tasks. Across domains, cognitive uncertainty always exhibits a pronounced hump shape. This means, first, that reported cognitive uncertainty is substantially higher for intermediate probabilities: our experimental participants tell us that they find it easier to think about lotteries with extreme payout probabilities, or about belief formation tasks that have extreme solutions. Second, measured cognitive uncertainty changes at a high rate close to zero and one but is much less sensitive for intermediate probabilities.

To investigate the implication of this non-linearity for observed experimental actions (certainty equivalents and stated beliefs), we conduct structural exercises that follow from the simple model in Section 2. Recall from equations (4) and (7) that a decision-maker's expected action (= implied subjective probability) can be expressed as a linear combination of p and the cognitive default p^d, where the relative weights depend on the magnitude of cognitive uncertainty. Again assuming $B = 1$ and setting $p^d = 1/2$:

$$\mathbb{E}[a^o] \equiv w(p) = (1 - \sigma_{CU}^2 \cdot D) \cdot p + \sigma_{CU}^2 \cdot D \cdot \frac{1}{2}, \tag{8}$$

where $D \equiv 1/\sigma_p^2$ is a parameter that linearly scales the square of measured cognitive uncertainty σ_{CU} into an action. The slope of this implied "weighting function" with respect to objective probabilities p is given by

$$\frac{\partial w(p)}{\partial p} = 1 - \sigma_{CU}^2 \cdot D + \frac{\partial \sigma_{CU}^2}{\partial p} \cdot D \cdot \left(\frac{1}{2} - p\right). \tag{9}$$

To illustrate the implication of this, consider the case of $p < 1/2$. Equation (9) says that the implied weighting function will be steeper if (i) cognitive uncertainty is lower and (ii) the sensitivity of cognitive uncertainty with respect to p is higher. Note that empirically both of these effects point in the direction that *the implied weighting function is steeper close to the extremes*.

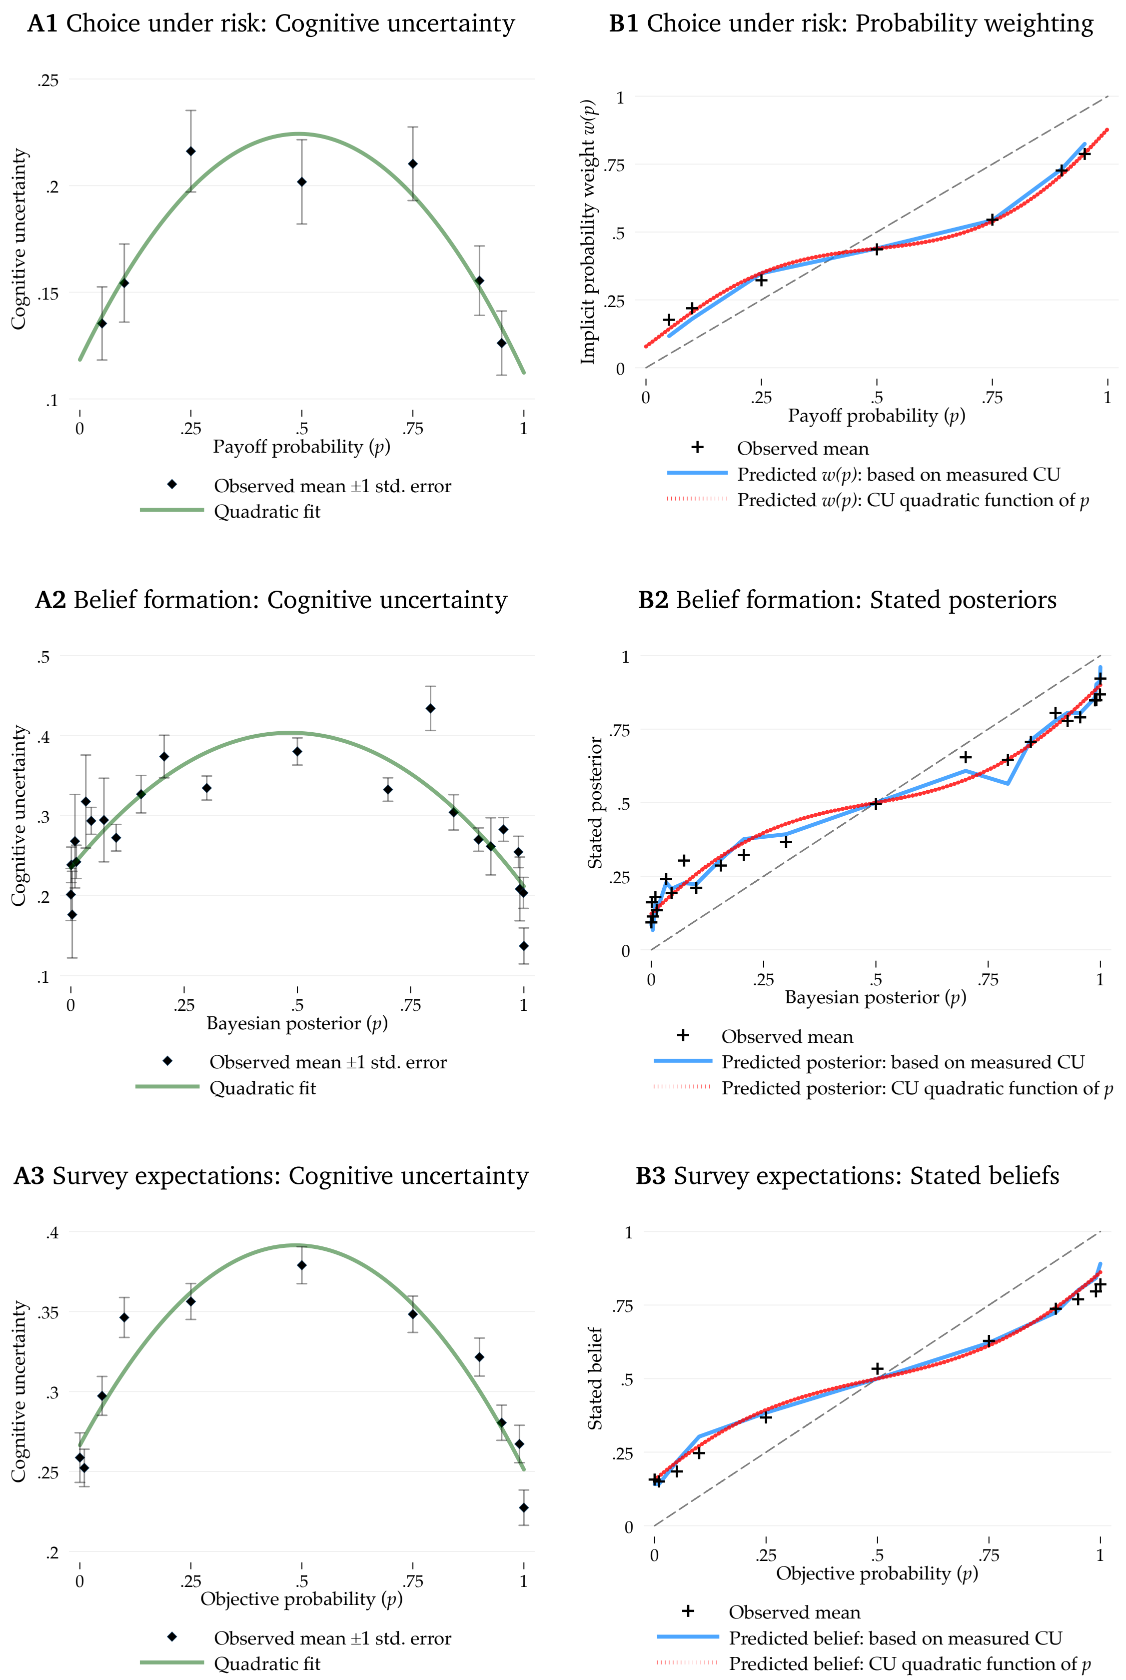

Figure 9: The observed non-linear relationship between measured cognitive uncertainty and probabilities (Panels A1-A3) generates the inverse-S shaped patterns in predicted choices and beliefs (Panels B1-B3) in a model of Bayesian shrinkage that is linear in p (see equations (4) and (8)). The samples include the same observations as in the baseline analyses in Sections 3–5.

To investigate whether the variation in measured cognitive uncertainty across objective probabilities could indeed produce the observed inverse S-shapes in certainty equivalents and stated beliefs, we estimate equation (8) using standard non-linear least squares techniques, and then plot predicted actions for each value of p. Recall that all quantities in equation (8) except for the scaling parameter D are observed. In choice under risk, we allow for potential curvature of the utility function by estimating both $\hat{D}$ and a CRRA parameter.[19] Unlike in other recent work that structurally estimates behavioral models, our primary interest is not so much in the estimated (nuisance) parameters per se but rather in how well the actions that are predicted from measured cognitive uncertainty fit the data.

Four remarks about this exercise are in order. First, conditional on σ_{CU}, the implied "weighting function" in equation (8) is *linear in p*. Thus, all non-linearities in predicted values must stem from the empirical dependence of σ_{CU} on p. Second, because self-reported cognitive uncertainty is likely measured with error, we estimate (8) at the population level after averaging a and σ_{CU} for each value of p. Third, our estimation takes participants' reported cognitive uncertainty at face value, up to a multiplicative constant D. In other words, we treat the data as if the self-reports had cardinal meaning, up to a scaling parameter. This scaling parameter is necessary both because we don't observe subjects' prior variance σ_p^2 and because our cognitive uncertainty elicitation does not specify whether participants are supposed to indicate, e.g., 95% or 99% confidence intervals.[20] Fourth, it is worth highlighting that our procedure does not mechanically generate good predictions. A good fit would mechanically arise if we used subjects' actions (certainty equivalents and beliefs) to "reverse engineer" cognitive uncertainty σ_{CU}^2. Instead, we don't estimate σ_{CU}^2 from subjects' actions but directly use our cognitive uncertainty data. We only use subjects actions' to estimate one value for the nuisance parameter D that is constant across p.

The solid blue line in the right-hand panels of Figure 9 plots predicted actions, i.e., the predicted values after estimating the scaling parameter D. Across decision domains, we see non-linearities that visually fit the observed data well. Because equation (8) is in principle linear in p, the characteristic inverse S-shape is entirely generated by how p affects cognitive uncertainty.

To make this point even more explicit, Figure 9 also plots an additional set of predicted actions that do not rely on *measured* cognitive uncertainty but instead on values

[19]With x the lottery payout, the estimation equation becomes

$$a = (w(p)\cdot x^{\gamma})^{\frac{1}{\gamma}} = \left(\left[(1-\sigma_{CU}^2\cdot D)\cdot p + \sigma_{CU}^2\cdot D\cdot 1/2\right]\cdot x^{\gamma}\right)^{\frac{1}{\gamma}}. \tag{10}$$

The non-linearity that results from γ is orthogonal to the non-linearity of $w(p)$ that we are interested in.

[20]As discussed in Section 3.1.2 and shown in Appendix B, subjects don't appear to be able to distinguish between these different confidence intervals in any case.

that are *predicted* from the objective probability p. Motivated by the near-quadratic relationships in the left-hand panels of Figure 9, we first generate predicted values of cognitive uncertainty from regressing measured cognitive uncertainty on probabilities:

$$\sigma_{CU} = \alpha + \beta_1 \cdot p + \beta_2 \cdot p^2 + u. \tag{11}$$

Then, we re-estimate equation (8) using the predicted values $\hat{\sigma}_{CU}$ instead of measured cognitive uncertainty. Again, it is worth pointing out that this procedure uses subjects' actions only for the purposes of estimating the scaling parameter D (and the CRRA parameter in choice under risk), while the variation in σ^2_{CU} entirely stems from the estimated relationship between measured cognitive uncertainty and objective probabilities.

As we can see from the dotted red lines in Figure 9, this procedure again delivers fitted values that exhibit pronounced inverse S-shapes, purely as a result of the dependence of cognitive uncertainty on objective probabilities. In summary, these results document that variable cognitive noise plays a meaningful role in generating inverse S-shapes.

6.2 Log Coding and the Gonzalez-Wu Weighting Function

The previous analysis assumed that people perceive p linearly, so that any non-linearities in $w(p)$ had to result from a genuine non-linear relationship between p and cognitive noise that occurs outside of the model. A complementary approach to shed light on this non-linearity is to entertain the possibility that human brains do not encode probabilities linearly. In particular, some work in cognitive science (Zhang and Maloney, 2012; Zhang et al., 2020) has entertained the possibility that the brain encodes frequencies and probabilities into a quantity q in log odds space by applying

$$q = Q(p) = \ln \frac{p}{1-p}. \tag{12}$$

As shown in the psychological literature, log-odds models fit observed behavior well. In our context, the assumption of log coding is appealing for two reasons. First, cognitive noise in log odds space delivers the specific non-linear relationship between measured cognitive uncertainty and p that we document empirically. Second, as we derive in Appendix G.1, the combination of log coding and cognitive noise delivers the following weighting function for the median decision maker (also see Khaw et al., 2017)):

$$w(p)^{G-W} := \frac{\delta p^\lambda}{\delta p^\lambda + (1-p)^\lambda}, \tag{13}$$

where $\delta = \exp\left((1-\lambda)\ln\frac{p^d}{1-p^d}\right)$ and λ is the same Bayesian shrinkage factor as in Section 2. This formulation is instructive because it corresponds to – and endogenizes – the well-known two-parameter specification of a probability weighting function popularized by Gonzalez and Wu (1999) and going back at least to Goldstein and Einhorn (1987).

In Appendix G.2, we estimate the parameters of this weighting function across decision domains. Comparing decisions that are associated with below (L) and above average (H) cognitive uncertainty, we estimate that the sensitivity parameter $\hat{\lambda}$ is 64% higher in choice under risk ($\hat{\lambda}_L = 0.54$ vs. $\hat{\lambda}_H = 0.33$). Likewise, the sensitivity parameter is 57% higher in belief updating ($\hat{\lambda}_L = 0.60$ vs. $\hat{\lambda}_H = 0.38$) and 146% higher in the pooled survey beliefs data ($\hat{\lambda}_L = 0.59$ vs. $\hat{\lambda}_H = 0.24$).

7 Subject-Level Heterogeneity in Cognitive Uncertainty

Next, we characterize heterogeneity in cognitive uncertainty across individuals. A natural starting point is to decompose the variation in our experimental data into between- and within-subject variation. For this purpose, we regress the collection of cognitive uncertainty statements on subject fixed effects, separately for each decision domain. We find that the variance explained is 44% in choice under risk, 53% in belief updating, and 60% in survey expectations. These numbers represent weak lower bounds for the fraction of the true variation that is due to between-subject variation, as they do not account for measurement error in cognitive uncertainty.

An additional way to investigate the existence of types is to look at subjects' consistency in cognitive uncertainty *across decision domains*. Recall that each subject completed the survey expectations tasks and additionally either the risky choice or the belief updating experiments. The correlation between average subject-level cognitive uncertainty in belief updating and average cognitive uncertainty in survey expectations is $\rho = 0.57$. The correlation between cognitive uncertainty in risky choice and survey expectations is $\rho = 0.35$. We conclude from these analyses that cognitive uncertainty varies in meaningful and reasonably stable ways across participants.

Figure 10 shows correlates of this individual-level variation. We report standardized beta coefficients, so that the y-axis shows the percent change in cognitive uncertainty that is associated with a 1% increase in the explanatory variable of interest. While the results are mixed overall, the strongest and most consistent correlations reflect that women, people who take less time to complete the task, and people with lower cognitive skills report higher cognitive uncertainty. The correlational evidence on response times and cognitive skills suggests that the availability of cognitive resources may reduce cognitive uncertainty. This is consistent with the view that cognitive noise is reduced as

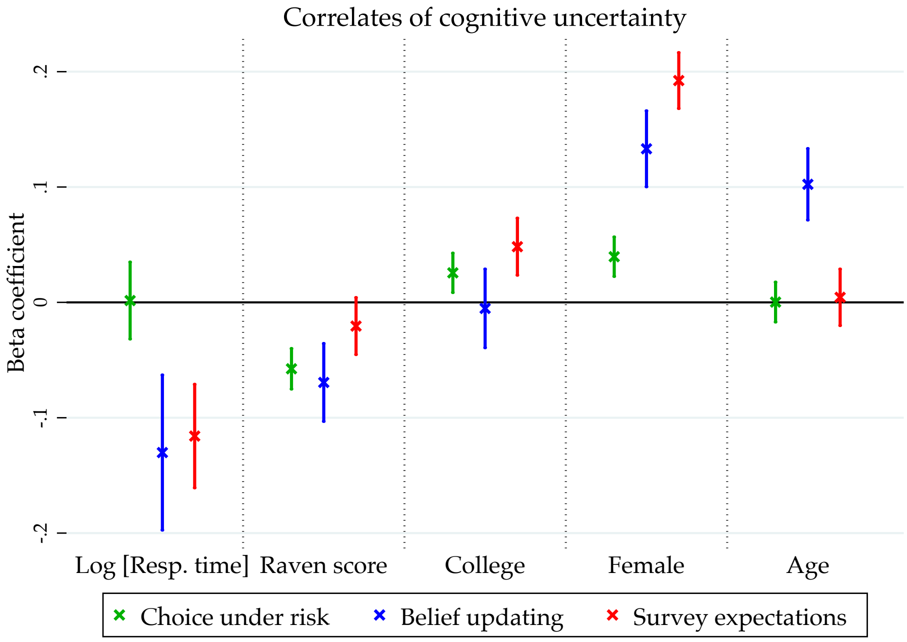

Figure 10: Correlates of average cognitive uncertainty. The figure shows the standardized beta coefficients of regressions of a subject's average cognitive uncertainty on different variables, controlling for treatment fixed effects. Error bars denote +/− 1 s.e.m. The values on the y-axis show the percent change in cognitive uncertainty that is associated with a 1% increase in the explanatory variable of interest. The beta coefficients are estimated conditional on treatment fixed effects. Response times are computed as total completion time within the relevant part of the experiment. $N = 1,000$ observations for choice under risk and belief updating and $N = 2,000$ observations for survey expectations.

participants sequentially accumulate "evidence" about a decision problem, as in sequential sampling and drift diffusion models (Krajbich et al., 2010; Fudenberg et al., 2018).[21]

8 Robustness and Discussion of Potential Confounds

Additional pre-registered analyses. The pre-registration specified that we conduct our analyses on three different samples: (i) excluding extreme outliers, as done thus far; (ii) using all data; and (iii) excluding "speeders," defined as subjects in the bottom decile of the response time distribution. Appendices C.2 and C.3 reproduce the analysis of choice under risk on the full sample and excluding speeders. Appendices D.5 and D.6 provide analogous analyses for belief updating and Appendices E.4 and E.5 for survey expectations. The results are always very similar.

[21]At the same time, Appendix H reports on a set of experiments that suggest that a moderate increase in incentives leads to longer response times but reduces cognitive uncertainty only marginally or not at all.

Random choice. Different literatures have proposed that the inverse S-shaped response functions that motivate our study are mechanically generated by random choice in combination with the fact that the response scale in choice and belief updating experiments is typically bounded by $p \in [0, 1]$.[22] While our perspective is perhaps seemingly similar in that it highlights the role of noise, we argue and document empirically that compressed responses reflect the *response to cognitive noise*, rather than decision noise itself. In particular, in all of our experimental contexts, random choice accounts always predict compression towards 50%, whereas we have shown that people's responses are generically compressed towards a cognitive default, even when it is different from 50%. Indeed, we have shown that decisions that are associated with higher cognitive uncertainty respond *more strongly* to manipulations of the cognitive default. This is directly predicted by our cognitive shrinkage model, but not by random choice models.[23]

Censoring. A related potential concern is that censoring generates some of our results. Censoring could either appear (i) in the action space (we did not allow participants to state probabilities above 100%) and / or (ii) in the cognitive uncertainty elicitation. As we discuss in detail in Appendix I, there is little indication of censoring in our data, and our results are always robust to excluding seemingly-censored observations.

Cognitive ability. It is well-known that the presence of behavioral anomalies tends to be correlated with low cognitive ability, including in the domains of belief updating (e.g., Hoppe and Kusterer, 2011; Enke and Zimmermann, 2019) and probability weighting (e.g., Choi et al., 2018). At the same time, the general notion that "low cognitive ability generates biases" has been relatively atheoretical and unspecific in that it does not explain how and why low cognitive ability should produce a specific behavioral pattern rather than just unsystematic noise. As discussed in Sections 2 and 7, our view is that cognitive noise is plausibly reduced through deliberation of the problem, such that higher cognitive skills (increased ability to deliberate) reduce cognitive noise. We therefore embrace the idea that our correlational results on cognitive uncertainty partly reflect heterogeneity in cognitive ability.

[22]See e.g. Blavatskyy (2007) and Blavatskyy and Pogrebna (2010) on decision noise in choice under risk and Erev et al. (1994), Moore and Healy (2008) and Costello and Watts (2014) on random errors in belief formation.

[23]An additional problem with a random choice explanation in our data is that the random choice component would have to be *very* large. While it may be plausible that – given decision bounds of [0,1] – random choice generates biased beliefs when the Bayesian posterior is close to one of the boundaries, it is considerably less plausible that subjects asymmetrically "hit the boundary" when the Bayesian posterior is far away from either boundary, such as for Bayesian posteriors of 30% or 70%. This is even more true given that we observe almost no censoring in our data. See also the discussion in Appendix I for additional evidence in line with shrinkage but not random choice.

At the same time, our paper provides ample evidence that the magnitude of cognitive noise does not just depend on cognitive ability but also on features of the decision problem: (i) our correlational results on choice under risk and belief updating remain statistically highly significant when we only consider within-subject-across-task variation in cognitive uncertainty and choices, eliminating the potential role of cognitive ability; (ii) our compound manipulations exogenously vary cognitive uncertainty within subject by making the decision problem more complex; and (iii) as documented in Section 6, cognitive uncertainty varies systematically over the probability range, which is crucial in generating inverse S-shaped responses, but cannot be explained by subject-level variation in cognitive ability. Finally, an atheoretical account of low cognitive ability does not explain the results of our default manipulations, while they are directly predicted by our model.

9 Discussion and Related Literature

This paper has shown that cognitive uncertainty predicts economic beliefs and actions, and that it allows us to bring together empirical regularities from decision tasks on choice under risk and ambiguity, belief updating, and survey expectations. Across all of these perhaps seemingly-unrelated decision domains, cognitive uncertainty strongly predicts behavior, which allows us to tie together regularities including the probability weighting function, the fourfold pattern of risk attitudes, ambiguity-insensitivity, base rate insensitivity, conservatism, sample proportion effects, and predictable overoptimism and -pessimism in economic forecasts.

Our work builds on the rich literatures that have studied probability weighting, ambiguity-insensitivity, belief updating, and survey expectations. We view as one of our primary contributions that we bring these voluminous literatures – that have hitherto evolved in isolation from each other – together under a common umbrella. The literatures on choice under risk and ambiguity have long speculated that insensitivity to probabilities is due to cognitive limitations (Viscusi, 1989; Wakker, 2010; Baillon et al., 2018). Similarly, the lab belief updating literature reviewed by Benjamin (2019) and the survey expectations literature have long posited the idea that people avoid stating extreme beliefs, and there is much evidence suggesting that 50:50 responses in surveys reflect some version of "not knowing the answer" (Fischhoff and Bruine De Bruin, 1999).

Our paper endogenizes these and related ideas through cognitive noise and resulting cognitive shrinkage processes. In doing so, our paper builds on a vibrant literature – first in cognitive science and increasingly also in economics – on Bayesian noisy cognition. Within economics, this literature is almost entirely theoretical in nature; the main idea behind these contributions is usually that noisy cognition induces agents to perform an implicit Bayesian update about objectively known quantities, see Chater et al. (2008), Woodford (2012, 2019), Khaw et al. (2017), Gabaix and Laibson (2017), Gabaix (2014, 2019), Frydman and Jin (2019), and Steiner and Stewart (2016).[24]

Because we define and empirically measure cognitive uncertainty as *awareness* about cognitive noise, our paper also connects to different lines of empirical work that have entertained the possibility that people may not always perfectly know what to do. For instance, research in psychology and neuroscience under the umbrella of "decision confidence" shows that confident subjects exhibit lower across-task variability in actions (e.g., De Martino et al., 2013, 2017; Polania et al., 2019) and that ratings or probability assessments typically exhibit a quadratic relationship with reported confidence (Lebreton et al., 2015). The idea of measuring different types of cognitive imprecision or randomness in choice is also increasingly gaining traction in the economics literature, in both choice under risk (Butler and Loomes, 2007; Cubitt et al., 2015; Agranov and Ortoleva, 2020; Agranov et al., 2020) and survey expectations (Giustinelli et al., 2019; Drerup et al., 2017). Our approach contributes to these literatures by documenting that cognitive uncertainty predicts the direction and magnitude of bias in different choice and belief formation contexts that are typically studied by economists.

An attractive feature of studying (awareness of) cognitive noise is that it may hold the promise of tying together empirical regularities from across the social and cognitive sciences. While we and other economists typically focus on high-level economic decisions, it is becoming increasingly clear that the idea of Bayesian noisy cognition and resulting compression effects also explains a vast array of regularities from psychological tasks (e.g., Chater et al., 2008; Petzschner et al., 2015). For instance, in building on our work, Xiang et al. (2020) show empirically that awareness about cognitive noise is strongly predictive of well-known compression (central tendency) effects in various perceptual judgment tasks.

In light of this vibrant emerging line of work, we believe that measuring cognitive uncertainty is likely to be productive for understanding economic decision-making going forward. First, future research may measure cognitive uncertainty to explain economic behavior also outside of the domain of probabilities that we study in this paper. Second, future work my shed light on the psychological mechanisms that actually generate cognitive noise, and to which extent this is affected by (implicit) cost-benefit considerations as posited by the literature on rational inattention.

Despite the promise that the idea of Bayesian noisy cognition holds for connecting

[24]In the psychology literature, Fennell and Baddeley (2012) present a model in which the probability weighting function arises through a Bayesian updating process that is formally similar to our stylized framework, except that the authors focus on *objective* (rather than cognitive) uncertainty, such as when people do not perfectly know the objective payout probabilities of a lottery.

seemingly-unrelated literatures and empirical regularities, it is worth repeating here that the mathematical framework upon which our analysis is based is stylized in nature and does not feature the richness of domain-specific models. This means that our approach will likely miss some relevant insights that domain-specific models can capture. These two approaches – identifying common principles and getting the full picture within a given domain right – are arguably both important and complement one another.

References

Agranov, Marina and Pietro Ortoleva, "Stochastic choice and preferences for randomization," *Journal of Political Economy*, 2017, *125* (1), 40–68.

__ **and** __ , "Ranges of Preferences and Randomization," *Working Paper*, 2020.

__ **, Paul J Healy, and Kirby Nielsen**, "Non-Random Randomization," *Available at SSRN*, 2020.

Ambuehl, Sandro and Shengwu Li, "Belief updating and the demand for information," *Games and Economic Behavior*, 2018, *109*, 21–39.

Baillon, Aurélien, Zhenxing Huang, Asli Selim, and Peter P Wakker, "Measuring ambiguity attitudes for all (natural) events," *Econometrica*, 2018, *86* (5), 1839–1858.

Benartzi, Shlomo and Richard H Thaler, "Naive diversification strategies in defined contribution saving plans," *American economic review*, 2001, *91* (1), 79–98.

Benjamin, Daniel J, "Errors in Probabilistic Reasoning and Judgmental Biases," in "Handbook of Behavioral Economics" 2019.

Bernheim, B Douglas and Charles Sprenger, "Direct tests of cumulative prospect theory," *Available at SSRN 3350196*, 2019.

Blavatskyy, Pavlo R, "Stochastic expected utility theory," *Journal of Risk and Uncertainty*, 2007, *34* (3), 259–286.

__ **and Ganna Pogrebna**, "Models of stochastic choice and decision theories: Why both are important for analyzing decisions," *Journal of Applied Econometrics*, 2010, *25* (6), 963–986.

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper, "Risk and rationality: Uncovering heterogeneity in probability distortion," *Econometrica*, 2010, *78* (4), 1375–1412.

Butler, David J and Graham C Loomes, "Imprecision as an account of the preference reversal phenomenon," *American Economic Review*, 2007, *97* (1), 277–297.

Chater, Nick, Mike Oaksford et al., *The probabilistic mind: Prospects for Bayesian cognitive science*, OUP Oxford, 2008.

Choi, Syngjoo, Jeongbin Kim, Eungik Lee, Jungmin Lee et al., "Probability Weighting and Cognitive Ability," Technical Report 2018.

Costello, Fintan and Paul Watts, "Surprisingly rational: Probability theory plus noise explains biases in judgment.," *Psychological review*, 2014, *121* (3), 463.

Cubitt, Robin P, Daniel Navarro-Martinez, and Chris Starmer, "On preference imprecision," *Journal of Risk and Uncertainty*, 2015, *50* (1), 1–34.

DellaVigna, Stefano and Devin Pope, "Predicting experimental results: who knows what?," *Journal of Political Economy*, 2018, *126* (6), 2410–2456.

Dimmock, Stephen G, Roy Kouwenberg, and Peter P Wakker, "Ambiguity attitudes in a large representative sample," *Management Science*, 2015, *62* (5), 1363–1380.

Drerup, Tilman, Benjamin Enke, and Hans-Martin Von Gaudecker, "The precision of subjective data and the explanatory power of economic models," *Journal of Econometrics*, 2017, *200* (2), 378–389.

Enke, Benjamin and Florian Zimmermann, "Correlation Neglect in Belief Formation," *Review of Economic Studies*, 2019, *86* (1), 313–332.

_ , Ricardo Rodriguez-Padilla, and Florian Zimmermann, "Moral universalism: Measurement and economic relevance," 2019.

_ , Uri Gneezy, Brian J Hall, David Martin, Vadim Nelidov, Theo Offerman, and Jeroen van de Ven, "Cognitive Biases: Mistakes or Missing Stakes?," 2020.

Erev, Ido, Thomas S Wallsten, and David V Budescu, "Simultaneous over-and underconfidence: The role of error in judgment processes.," *Psychological review*, 1994, *101* (3), 519.

Falk, Armin and Florian Zimmermann, "Consistency as a Signal of Skills," *Management Science*, 2017, *63* (7), 2197–2210.

_ , Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global evidence on economic preferences," *The Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

Fennell, John and Roland Baddeley, "Uncertainty plus prior equals rational bias: An intuitive Bayesian probability weighting function.," *Psychological Review*, 2012, *119* (4), 878.

Fischhoff, Baruch and Wändi Bruine De Bruin, "Fifty–fifty= 50%?," *Journal of Behavioral Decision Making*, 1999, *12* (2), 149–163.

Fox, Craig R and Robert T Clemen, "Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior," *Management Science*, 2005, *51* (9), 1417–1432.

Frydman, Cary and Lawrence Jin, "Efficient Coding and Risky Choice," *Working Paper*, 2019.

Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki, "Speed, accuracy, and the optimal timing of choices," *American Economic Review*, 2018, *108* (12), 3651–84.

Gabaix, Xavier, "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

_ , "Behavioral inattention," in Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics-Foundations and Applications 2*, Elsevier, 2019, p. 261.

_ **and David Laibson**, "Myopia and discounting," Technical Report, National bureau of economic research 2017.

Gillen, Ben, Erik Snowberg, and Leeat Yariv, "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Journal of Political Economy*, 2019.

Giustinelli, Pamela, Charles F Manski, and Francesca Molinari, "Precise or Imprecise Probabilities? Evidence from Survey Response on Late-onset Dementia," Technical Report, National Bureau of Economic Research 2019.

Goldstein, William M and Hillel J Einhorn, "Expression theory and the preference reversal phenomena.," *Psychological review*, 1987, *94* (2), 236.

Gonzalez, Richard and George Wu, "On the shape of the probability weighting function," *Cognitive psychology*, 1999, *38* (1), 129–166.

Grether, David M., "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 1980, *95*, 537–557.

Halevy, Yoram, "Ellsberg revisited: An experimental study," *Econometrica*, 2007, *75* (2), 503–536.

Hoppe, Eva I and David J Kusterer, "Behavioral biases and cognitive reflection," *Economics Letters*, 2011, *110* (2), 97–100.

Hossain, Tanjim and Ryo Okui, "The binarized scoring rule," *Review of Economic Studies*, 2013, *80* (3), 984–1001.

Imas, Alex, Michael Kuhn, and Vera Mironova, "Waiting to choose," *Working Paper*, 2016.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 1979, *47* (2), 263–292.

Khaw, Mel Win, Ziang Li, and Michael Woodford, "Risk aversion as a perceptual bias," Technical Report, National Bureau of Economic Research 2017.

Krajbich, Ian, Carrie Armel, and Antonio Rangel, "Visual fixations and the computation and comparison of value in simple choice," *Nature neuroscience*, 2010, *13* (10), 1292.

Lebreton, Maël, Raphaëlle Abitbol, Jean Daunizeau, and Mathias Pessiglione, "Automatic integration of confidence in the brain valuation signal," *Nature neuroscience*, 2015, *18* (8), 1159.

Li, Chen, Uyanga Turmunkh, and Peter P Wakker, "Social and Strategic Ambiguity versus Betrayal Aversion," *Working Paper*, 2019.

Liang, Yucheng, "Learning from Unknown Information Sources," *Working Paper*, 2019.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa, "Failures in Contingent Reasoning: The Role of Uncertainty," *American Economic Review*, 2019.

Martino, Benedetto De, Sebastian Bobadilla-Suarez, Takao Nouguchi, Tali Sharot, and Bradley C Love, "Social information is integrated into value and confidence judgments according to its reliability," *Journal of Neuroscience*, 2017, *37* (25), 6066–6074.

_ , **Stephen M Fleming, Neil Garrett, and Raymond J Dolan**, "Confidence in value-based choice," *Nature neuroscience*, 2013, *16* (1), 105–110.

Moore, Don A and Paul J Healy, "The trouble with overconfidence.," *Psychological review*, 2008, *115* (2), 502.

Oprea, Ryan, "What Makes a Rule Complex?," *American Economic Review*, 2020, *110* (12), 3913–51.

Paolacci, Gabriele and Jesse Chandler, "Inside the Turk: Understanding Mechanical Turk as a participant pool," *Current Directions in Psychological Science*, 2014, *23* (3), 184–188.

Petzschner, Frederike H, Stefan Glasauer, and Klaas E Stephan, "A Bayesian perspective on magnitude estimation," *Trends in cognitive sciences*, 2015, *19* (5), 285–293.

Polania, Rafael, Michael Woodford, and Christian C Ruff, "Efficient coding of subjective value," *Nature neuroscience*, 2019, *22* (1), 134–142.

Stango, Victor and Jonathan Zinman, "We are all behavioral, more or less: A taxonomy of consumer decision making," *NBER Working Paper*, 2020, (w28138).

Starmer, Chris and Robert Sugden, "Testing for juxtaposition and event-splitting effects," *Journal of risk and uncertainty*, 1993, *6* (3), 235–254.

Steiner, Jakub and Colin Stewart, "Perceiving prospects properly," *American Economic Review*, 2016, *106* (7), 1601–31.

Trautmann, Stefan T and Gijs Van De Kuilen, "Ambiguity attitudes," *The Wiley Blackwell handbook of judgment and decision making*, 2015, *1*, 89–116.

Tversky, Amos and Daniel Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and uncertainty*, 1992, *5* (4), 297–323.

Viscusi, W Kip, "Prospective reference theory: Toward an explanation of the paradoxes," *Journal of risk and uncertainty*, 1989, *2* (3), 235–263.

Wakker, Peter P, *Prospect theory: For risk and ambiguity*, Cambridge university press, 2010.

Woodford, Michael, "Prospect theory as efficient perceptual distortion," *American Economic Review*, 2012, *102* (3), 41–46.

—, "Modeling Imprecision in Perception, Valuation and Choice," *Working Paper*, 2019.

Xiang, Yang, Thomas Graeber, Benjamin Enke, and Samuel J Gershman, "Bayesian signatures of confidence and central tendency in perceptual judgment," 2020.

Zhang, Hang and Laurence T Maloney, "Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition," *Frontiers in neuroscience*, 2012, *6*, 1.

—, **Xiangjuan Ren, and Laurence T Maloney**, "The bounded rationality of probability distortion," *BioRxiv*, 2020, p. 662429.

ONLINE APPENDIX

A Nonlinear Extension of Theoretical Framework

Complementing the basic framework presented in Section 2, we now allow the rational action to be a nonlinear function of p. Specifically, the agent's objective function is given by

$$\min_a \quad v(a,p) = \frac{1}{2}\left(a - f(p)\right)^2. \tag{14}$$

This formulation more realistically captures the key elements of choice under risk. In particular, for a neoclassical agent, the certainty equivalent for a p% chance of getting \$1 is given by $CE = u^{-1}(p)$, so that $f(\cdot) = u^{-1}(\cdot)$.

We make the following simplifying assumptions. First, following some of the previous literature (Gabaix, 2019, 2014), we assume that the agent still chooses an action based on the posterior expectation about p,

$$\hat{a}^{nl} = f(\mathbb{E}[p|s]), \tag{15}$$

which is the analogue of equation (4). Second, we assume that the function $f(\cdot)$ is strictly monotone, such we can use a standard feature of the median of strictly monotone nonlinear functions. We derive properties of the median action a^e,

$$a^e(p) = \text{Median}(\hat{a}^{nl}|p) = f(\lambda p + (1-\lambda)p^d). \tag{16}$$

To ease notation, we set $x = p - p^d$, i.e., x is defined as the deviation from the default. This merely allows us to suppress the default in the expression for cognitive uncertainty.

We define cognitive uncertainty analogously to the linear case as the agent's perceived uncertainty about his optimal action. Instead of stating the standard deviation of the subjective posterior distribution, we specify cognitive uncertainty as the interquartile range around the observed action (which corresponds to 1.35 standard deviations):

$$\sigma^{nl}_{CU}(x) = \left| f\left(\lambda x + \frac{1}{2}\sqrt{1-\lambda}\sigma_x\right) - f\left(\lambda x - \frac{1}{2}\sqrt{1-\lambda}\sigma_x\right)\right|. \tag{17}$$

At the median, using $a^e(x) = f(\lambda x)$ from equation (16 yields

$$\sigma^{nl}_{CU}(x) = \left| a^e\left(x + \frac{1}{2}\frac{\sqrt{1-\lambda}}{\lambda}\sigma_x\right) - a^e\left(x - \frac{1}{2}\frac{\sqrt{1-\lambda}}{\lambda}\sigma_x\right)\right|. \tag{18}$$

A Taylor expansion of (18) gives

$$\sigma_{CU}^{nl} = \left|a^{e\prime}(x)\right| \frac{\sqrt{1-\lambda}}{\lambda}\sigma_x, \tag{19}$$

which is the nonlinear equivalent of equation (6):

$$\frac{\lambda}{\sqrt{1-\lambda}} = \frac{|a^{e\prime}(x)|\,\sigma_x}{\sigma_{CU}^{nl}}. \tag{20}$$

This nonlinear extension therefore provides the same comparative static predictions as the linear version.

B Calibrating the Cognitive Uncertainty Measurement

In all of our experiments, the elicitation of cognitive uncertainty did not specify which particular version of a subjective confidence interval we intend to elicit, such as a 90%, 95%, 99% or 100% confidence interval. We deliberately designed our experiments in this fashion because the hypothesis that underlines our research is that people have a hard time translating "99% confidence" into a statement about e.g. their certainty equivalent. In an attempt to trade off subject comprehension and quantitative interpretation, we hence refrained from inducing a particular version of a confidence interval.

To provide evidence for our conjecture that respondents cannot really tell the difference between different types of confidence intervals, we implemented an additional set of choice under risk experiments in which we elicited different versions of subjective confidence intervals. In these experiments, subjects were specifically instructed to state an interval such that they are "y% certain" that to them the lottery is worth between a and b. Across experimental conditions, y varied from 75% to 90% to 95% to 99% to 100%. To analyze these data, we compare average cognitive uncertainty within a treatment with average cognitive uncertainty in our baseline treatments, in which we did not provide a specific quantitative version of a confidence interval. In total, we ran these experiments with $N = 293$ subjects.

Figure 11 summarizes the results. Here, we plot the coefficients of the different treatment dummies in a regression with stated cognitive uncertainty as dependent variable. In this regression, the omitted category is our (unspecific) baseline treatment. Each coefficient hence corresponds to the implied difference in cognitive uncertainty between a treatment and our baseline treatment. There are two main results. First, cognitive uncertainty does not vary in meaningful ways across conditions: subjects state statistically indistinguishable cognitive uncertainty intervals, regardless of whether we specify them as 75%, 90% etc. interval. Second, if anything, reported cognitive uncertainty is *higher* in the more precise quantitative versions relative to our baseline version, as can be inferred from the positive point estimates. This again suggests that subjects have a harder time thinking about specific quantitative versions of a confidence interval relative to our more intuitive question. We conclude from this exercise that a more precise quantitative implementation of our cognitive uncertainty interval is unlikely to deliver a more helpful quantitative interpretation of our measure.

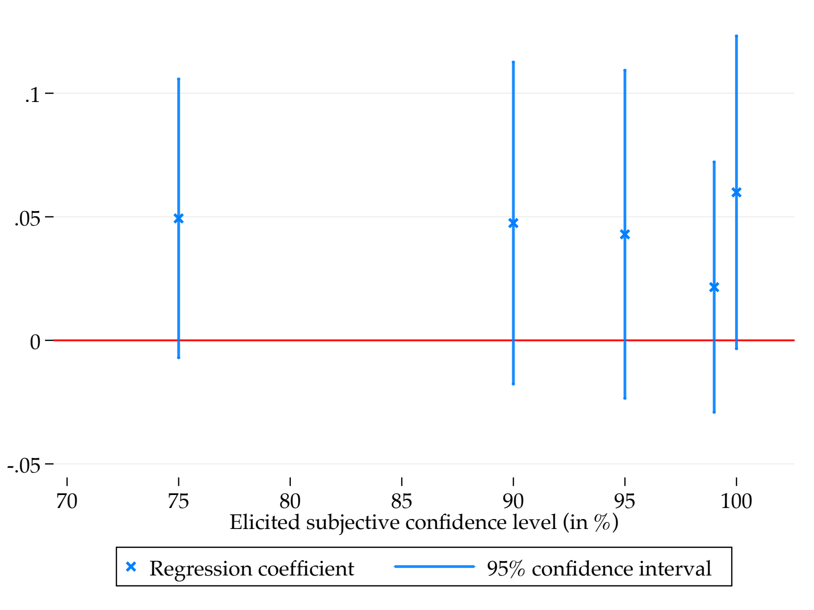

Figure 11: Comparison of average cognitive uncertainty across different elicitation modes in choice under risk. Each dot represents the coefficient of a treatment dummy in a regression with cognitive uncertainty as dependent variable. The explanatory variables are fixed effects for the different specifications of cognitive uncertainty, where the omitted category is our baseline wording. The plot controls for lottery amount fixed effects and probability of payout fixed effects.

C Additional Details and Analyses for Choice under Risk Experiments

C.1 Additional Figures

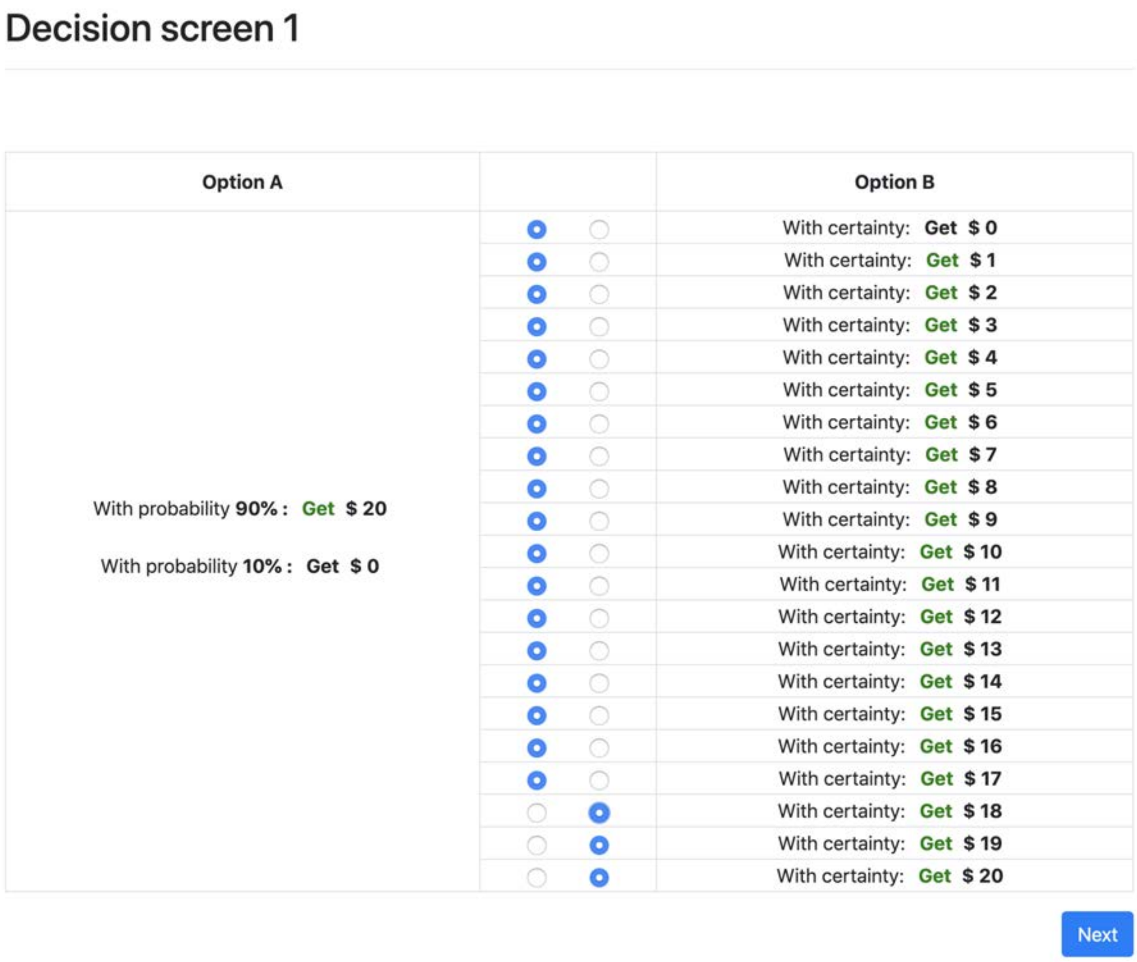

Figure 12: Decision screen to elicit certainty equivalents for lotteries

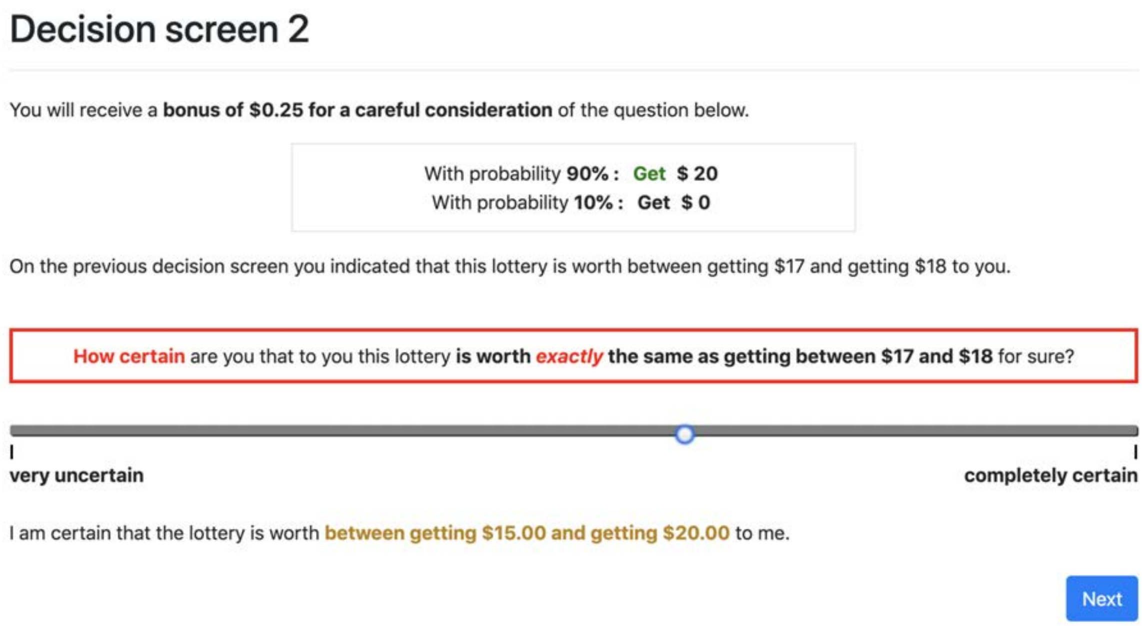

Figure 13: Decision screen to elicit cognitive uncertainty in choice under risk

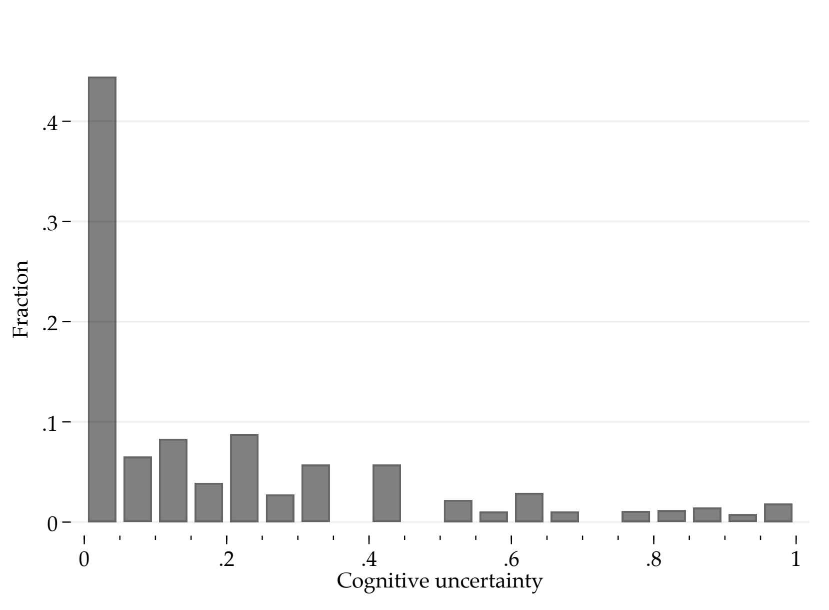

Figure 14: Histogram of cognitive uncertainty in baseline choice under risk tasks

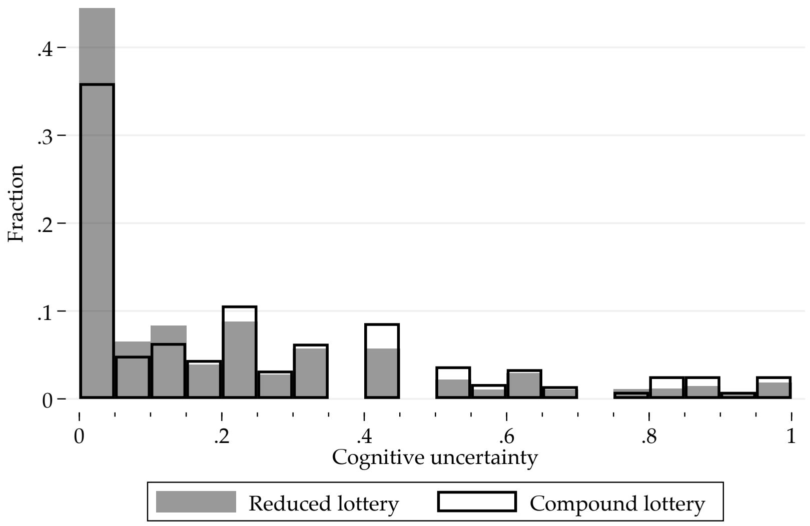

Figure 15: Histograms of cognitive uncertainty in choice under risk tasks, separately for reduced and compound lotteries

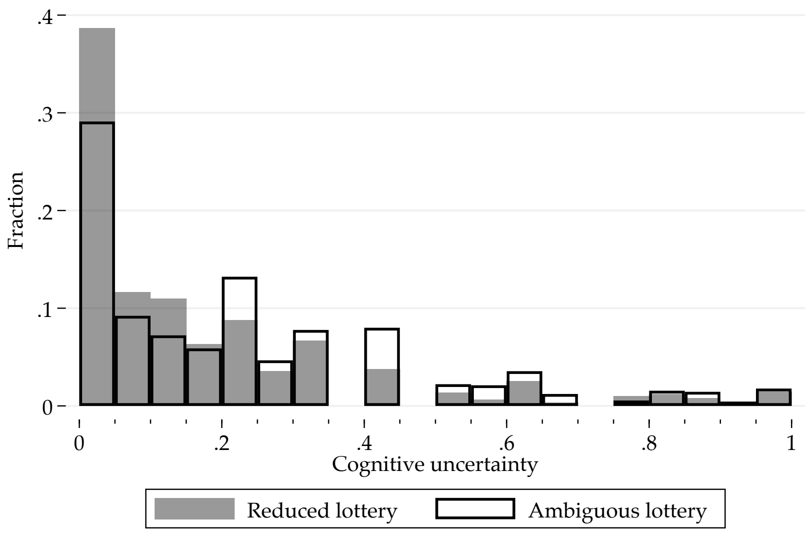

Figure 16: Histograms of cognitive uncertainty in choice under risk tasks, separately for reduced and ambiguous lotteries

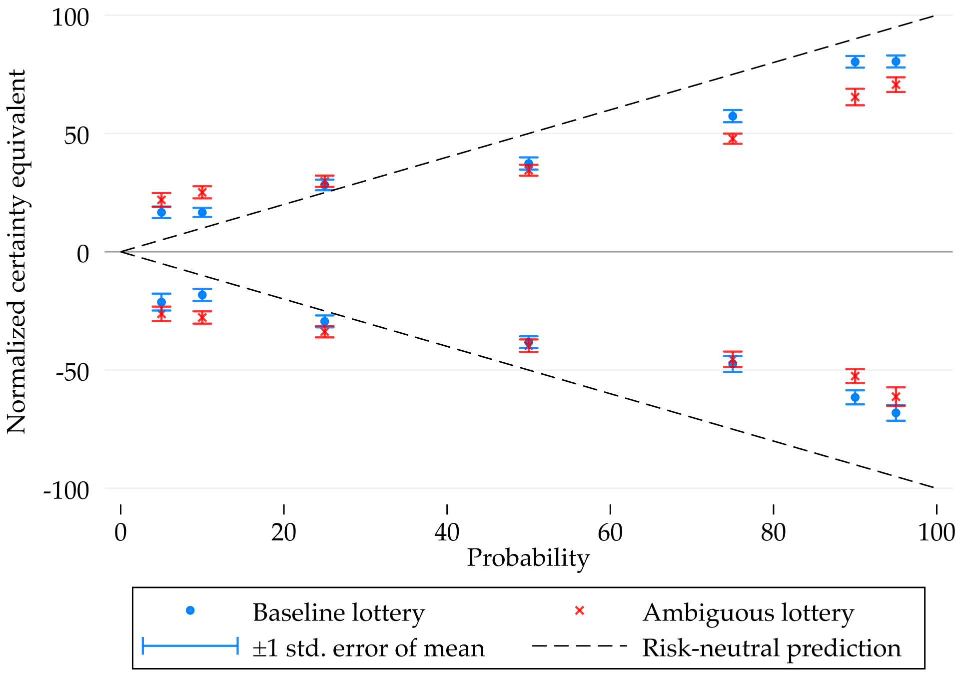

Figure 17: "Probability" weighting function separately for reduced and ambiguous lotteries. The payout "probability" for ambiguous lotteries is denoted by the midpoint of the interval of possible payout probabilities. The plot shows averages and corresponding standard error bars. The figure is based on 1,796 certainty equivalents of 300 subjects.

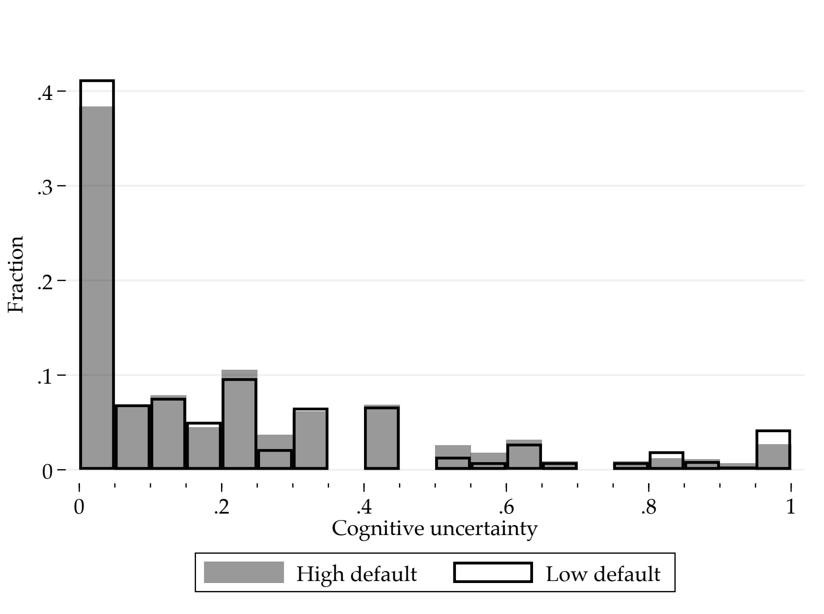

Figure 18: Histograms of cognitive uncertainty in choice under risk tasks, separately for treatments *High Default Risk* and *Low Default Risk*.

C.2 Results with Full Sample

Table 5: Insensitivity to probability and cognitive uncertainty (full sample)

	Dependent variable: Absolute normalized certainty equivalent					
	Gains		Losses		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Probability of payout	0.67***	0.67***	0.46***	0.46***	0.57***	0.57***
	(0.02)	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)
Probability of payout × Cognitive uncertainty	-0.51***	-0.51***	-0.087	-0.073	-0.29***	-0.28***
	(0.10)	(0.10)	(0.09)	(0.09)	(0.07)	(0.07)
Cognitive uncertainty	16.1***	16.1***	13.7**	13.2**	14.9***	15.4***
	(5.77)	(5.82)	(5.34)	(5.30)	(4.09)	(4.13)
Session FE	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes
Observations	1286	1286	1315	1315	2601	2601
R^2	0.49	0.50	0.27	0.29	0.36	0.36

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's normalized certainty equivalent, computed as midpoint of the switching interval divided by the non-zero payout. The sample includes choices from all baseline gambles with strictly interior probabilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Choice under risk: Baseline versus compound / ambiguous lotteries (full sample)

	Dependent variable: Absolute normalized certainty equivalent					
	Risk vs. compound risk			Risk vs. ambiguity		
	Gains	Losses	Pooled	Gains	Losses	Pooled
	(1)	(2)	(3)	(4)	(5)	(6)
Probability of payout	0.59***	0.45***	0.57***	0.72***	0.51***	0.68***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.04)	(0.03)
Probability of payout × 1 if compound / ambiguous lottery	-0.34***	-0.25***	-0.28***	-0.22***	-0.17***	-0.17***
	(0.04)	(0.04)	(0.03)	(0.03)	(0.04)	(0.02)
Probability of payout × Cognitive uncertainty			-0.25***			-0.45***
			(0.05)			(0.10)
1 if compound / ambiguous lottery	13.6***	12.3***	12.2***	8.02***	9.09***	7.06***
	(2.09)	(1.98)	(1.44)	(1.15)	(2.40)	(1.33)
Cognitive uncertainty			13.2***			20.7***
			(3.42)			(6.16)
Session FE	No	No	Yes	No	No	Yes
Demographic controls	No	No	Yes	No	No	Yes
Observations	1958	1947	3905	900	900	1800
R^2	0.37	0.21	0.29	0.52	0.27	0.42

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's absolute normalized certainty equivalent. In columns (1)–(3), the sample includes choices from the baseline and compound lotteries, where for comparability the set of baseline lotteries is restricted to lotteries with payout probabilities of 10%, 25%, 50%, 75%, and 90%, see Figure 3. In columns (4)–(6), the sample includes choices from the baseline and ambiguous lotteries. For ambiguous lotteries, we define the payout "probability" as the midpoint of the interval of possible payout probabilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Choice under risk: Treatments *Low Default* and *High Default* (full sample)

	Dependent variable: Absolute normalized certainty equivalent					
	Gains		Losses		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
0 if *High Default*, 1 if *Low Default*	-11.9***	-11.4***	-3.28	-2.82	-7.61***	-7.08***
	(1.93)	(1.98)	(2.28)	(2.25)	(1.60)	(1.60)
Probability of payout	0.56***	0.56***	0.51***	0.51***	0.54***	0.54***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)
Probability of payout × Cognitive uncertainty	-0.49***	-0.50***	-0.26**	-0.28**	-0.36***	-0.38***
	(0.10)	(0.10)	(0.13)	(0.13)	(0.09)	(0.09)
Cognitive uncertainty	16.7***	16.9***	22.3***	22.5***	19.1***	19.4***
	(6.08)	(6.05)	(8.21)	(8.40)	(5.19)	(5.25)
Session FE	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes
Observations	900	900	900	900	1800	1800
R^2	0.35	0.36	0.23	0.26	0.27	0.29

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's normalized certainty equivalent, computed as midpoint of the switching interval divided by the non-zero payout. The sample includes choices from treatments *Low Default* and *High Default*. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.3 Results excluding Speeders

Table 8: Insensitivity to probability and cognitive uncertainty (excl. speeders)

	Dependent variable: Absolute normalized certainty equivalent					
	Gains		Losses		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Probability of payout	0.68***	0.68***	0.47***	0.47***	0.58***	0.58***
	(0.02)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)
Probability of payout × Cognitive uncertainty	-0.57***	-0.56***	-0.12	-0.10	-0.33***	-0.32***
	(0.10)	(0.10)	(0.09)	(0.09)	(0.08)	(0.08)
Cognitive uncertainty	19.0***	19.0***	15.1***	14.4**	17.0***	17.3***
	(6.29)	(6.26)	(5.62)	(5.59)	(4.38)	(4.41)
Session FE	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes
Observations	1162	1162	1187	1187	2349	2349
R^2	0.49	0.50	0.27	0.29	0.36	0.36

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's normalized certainty equivalent, computed as midpoint of the switching interval divided by the non-zero payout. The sample includes choices from all baseline gambles with strictly interior probabilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Choice under risk: Baseline versus compound lotteries (excl. speeders)

	Dependent variable: Absolute normalized certainty equivalent					
	Risk vs. compound risk			Risk vs. ambiguity		
	Gains	Losses	Pooled	Gains	Losses	Pooled
	(1)	(2)	(3)	(4)	(5)	(6)
Probability of payout	0.59***	0.46***	0.57***	0.73***	0.51***	0.68***
	(0.02)	(0.03)	(0.02)	(0.03)	(0.04)	(0.03)
Probability of payout × 1 if compound / ambiguous lottery	-0.32***	-0.23***	-0.25***	-0.24***	-0.18***	-0.17***
	(0.04)	(0.04)	(0.03)	(0.03)	(0.04)	(0.03)
Probability of payout × Cognitive uncertainty			-0.26***			-0.48***
			(0.06)			(0.10)
1 if compound / ambiguous lottery	12.5***	11.6***	11.2***	8.52***	9.85***	7.07***
	(2.18)	(2.05)	(1.51)	(1.29)	(2.53)	(1.46)
Cognitive uncertainty			14.2***			27.4***
			(3.70)			(5.84)
Session FE	No	No	Yes	No	No	Yes
Demographic controls	No	No	Yes	No	No	Yes
Observations	1766	1753	3519	774	834	1608
R^2	0.38	0.22	0.30	0.54	0.27	0.42

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's absolute normalized certainty equivalent. In columns (1)–(3), the sample includes choices from the baseline and compound lotteries, where for comparability the set of baseline lotteries is restricted to lotteries with payout probabilities of 10%, 25%, 50%, 75%, and 90%, see Figure 3. In columns (4)–(6), the sample includes choices from the baseline and ambiguous lotteries. For ambiguous lotteries, we define the payout "probability" as the midpoint of the interval of possible payout probabilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Choice under risk: Treatments *Low Default* and *High Default* (excl. speeders)

	Dependent variable: Absolute normalized certainty equivalent					
	Gains		Losses		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
0 if *High Default*, 1 if *Low Default*	-11.9***	-11.4***	-3.28	-2.82	-7.61***	-7.08***
	(1.93)	(1.98)	(2.28)	(2.25)	(1.60)	(1.60)
Probability of payout	0.56***	0.56***	0.51***	0.51***	0.54***	0.54***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)
Probability of payout × Cognitive uncertainty	-0.49***	-0.50***	-0.26**	-0.28**	-0.36***	-0.38***
	(0.10)	(0.10)	(0.13)	(0.13)	(0.09)	(0.09)
Cognitive uncertainty	16.7***	16.9***	22.3***	22.5***	19.1***	19.4***
	(6.08)	(6.05)	(8.21)	(8.40)	(5.19)	(5.25)
Session FE	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes
Observations	900	900	900	900	1800	1800
R^2	0.35	0.36	0.23	0.26	0.27	0.29

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's normalized certainty equivalent, computed as midpoint of the switching interval divided by the non-zero payout. The sample includes choices from treatments *Low Default* and *High Default*. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Additional Details and Analyses for Belief Updating Experiments

D.1 Additional Figures

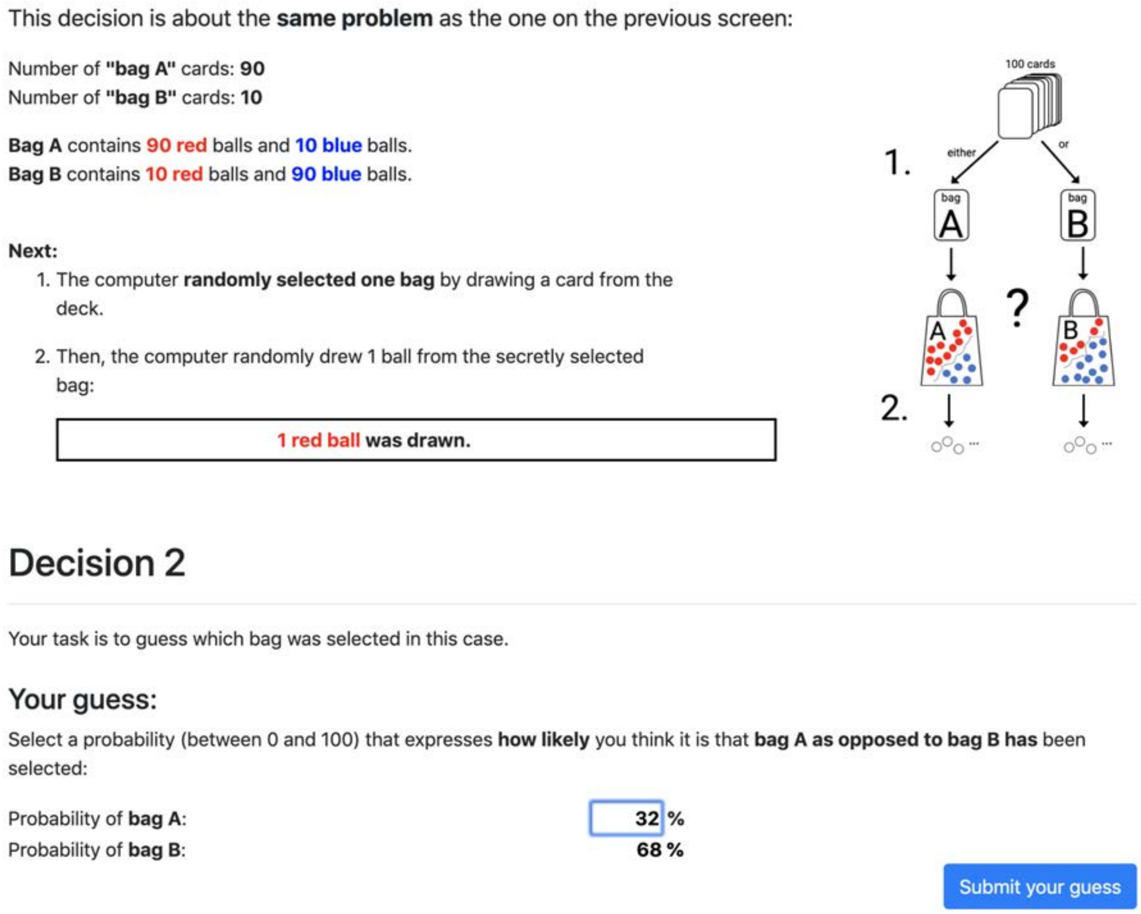

Figure 19: Decision screen to elicit posterior belief in belief updating tasks

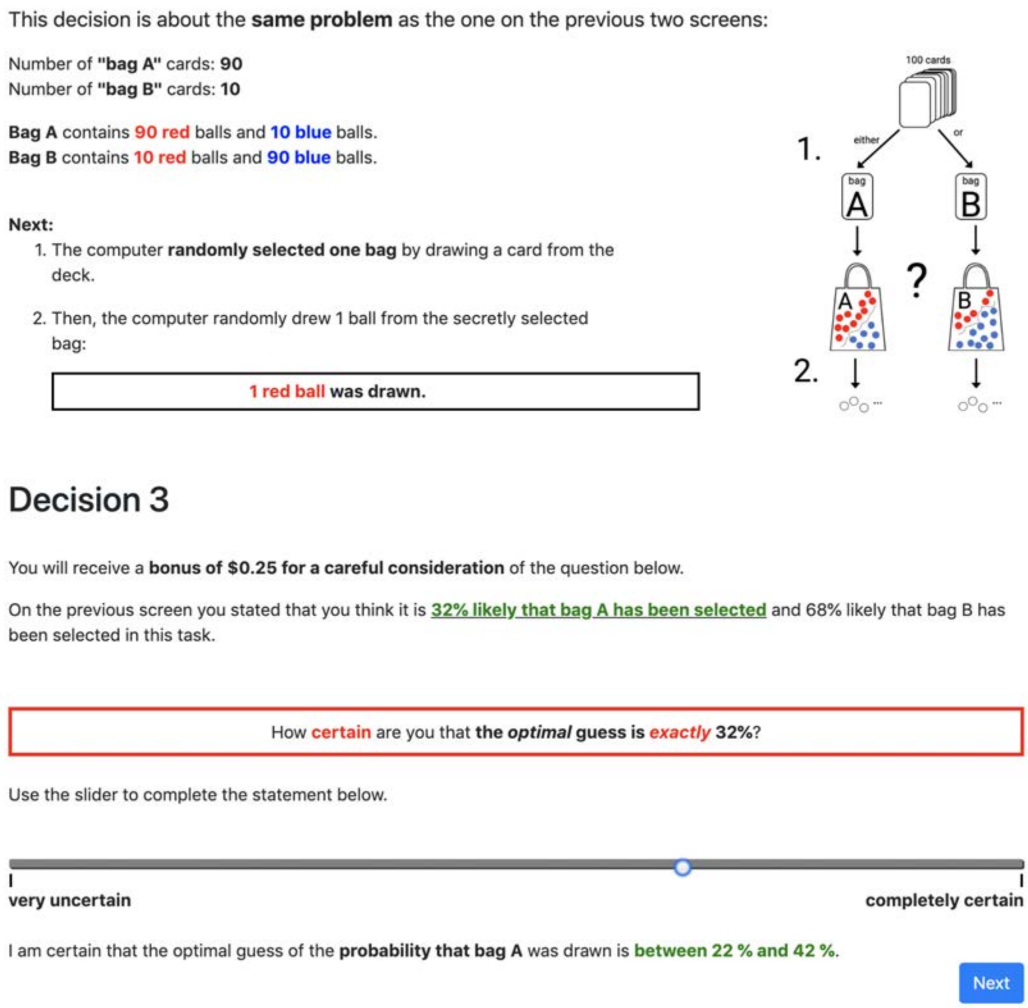

Figure 20: Decision screen to elicit cognitive uncertainty in belief updating

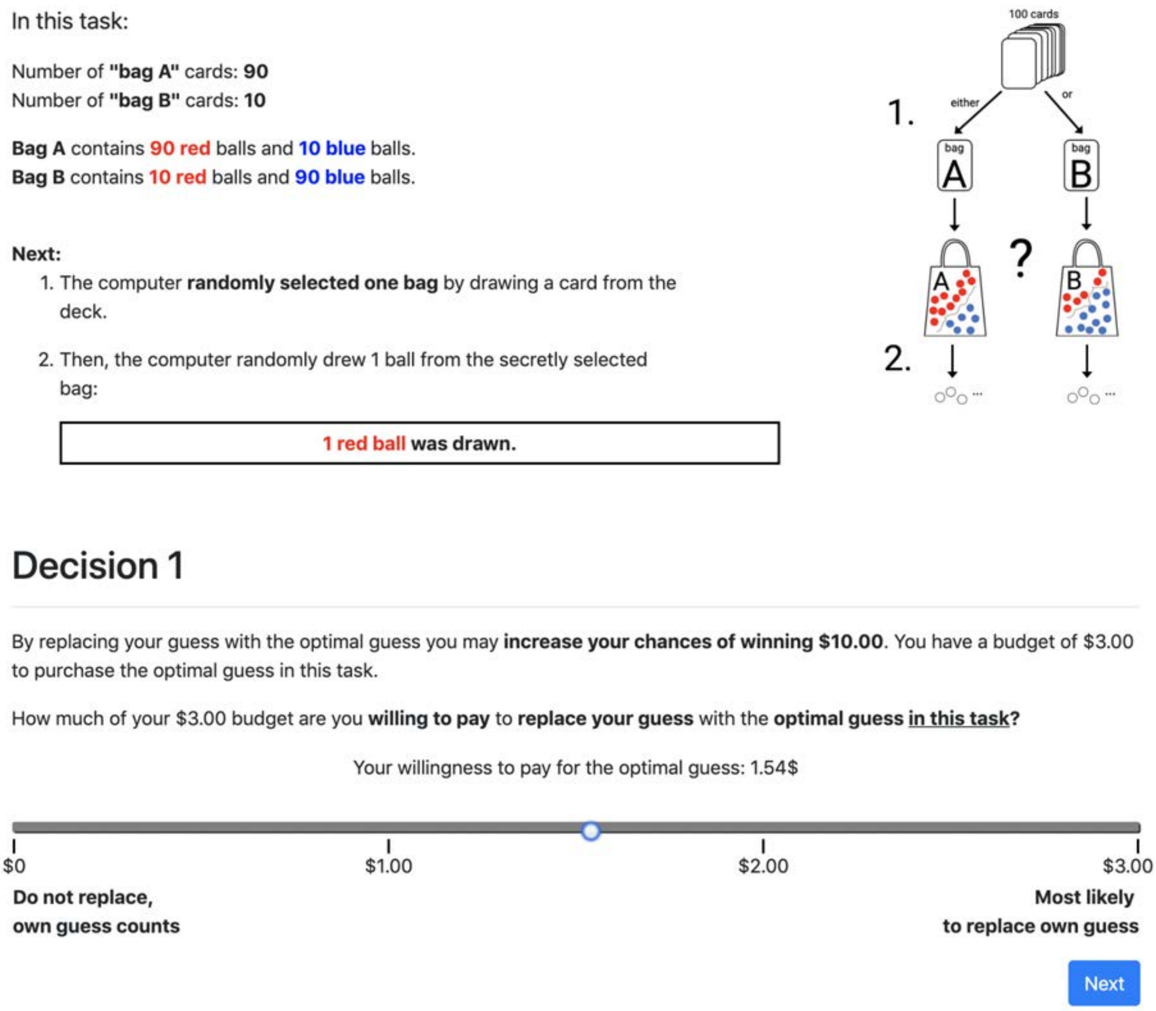

Figure 21: Decision screen to elicit willingness-to-pay for optimal guess in belief updating

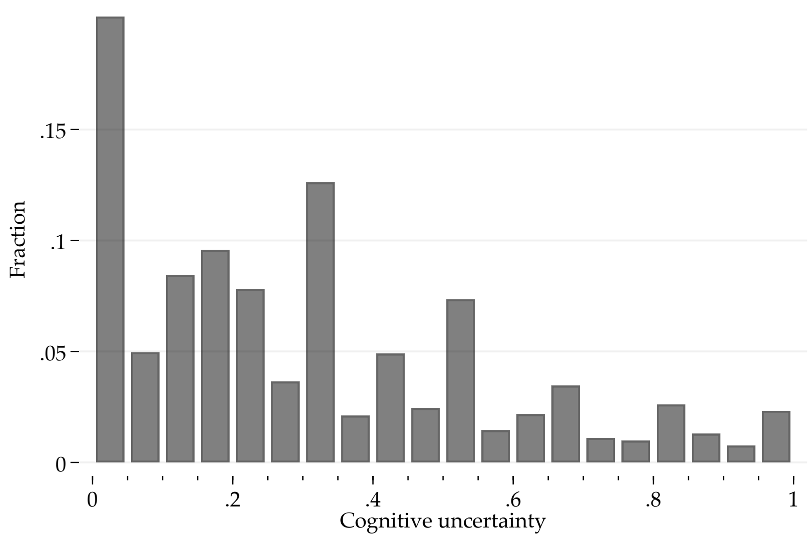

Figure 22: Histogram of cognitive uncertainty in baseline belief updating tasks

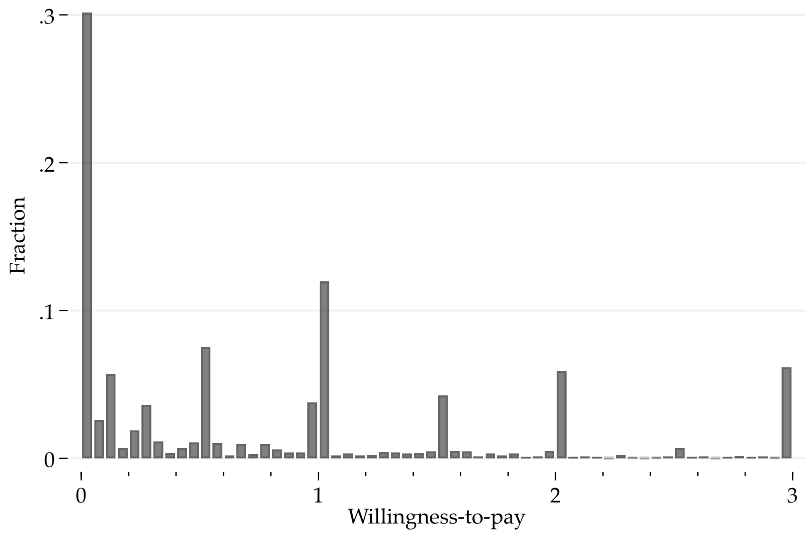

Figure 23: Histogram of willingness-to-pay to replace own guess by Bayesian posterior in baseline belief updating tasks

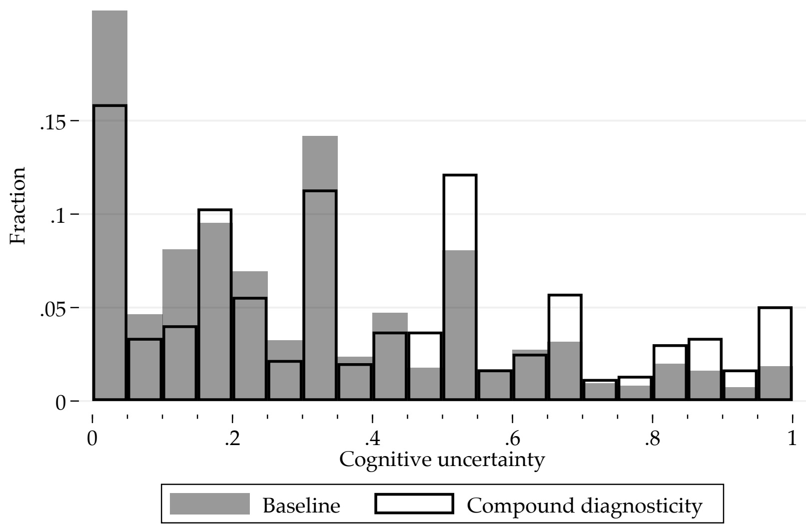

Figure 24: Histograms of cognitive uncertainty in belief updating tasks, separately for baseline and compound diagnosticities

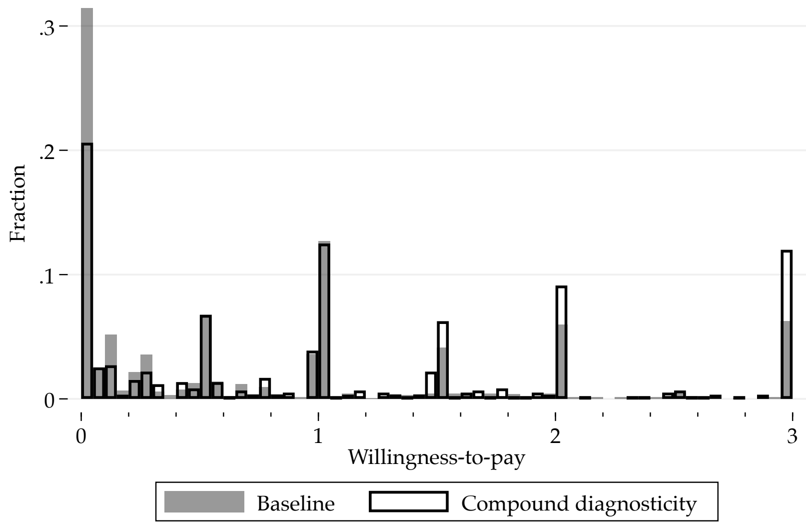

Figure 25: Histograms of willingness-to-pay to replace own guess by Bayesian posterior in belief updating tasks, separately for baseline and compound diagnosticities

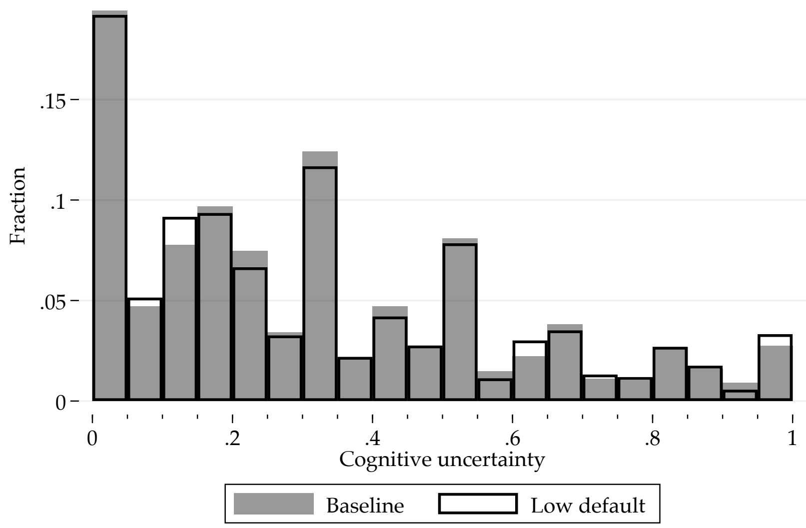

Figure 26: Histograms of cognitive uncertainty in belief updating tasks, separately for treatments *Baseline* and *Low Default Beliefs*.

D.2 Results with WTP Measure

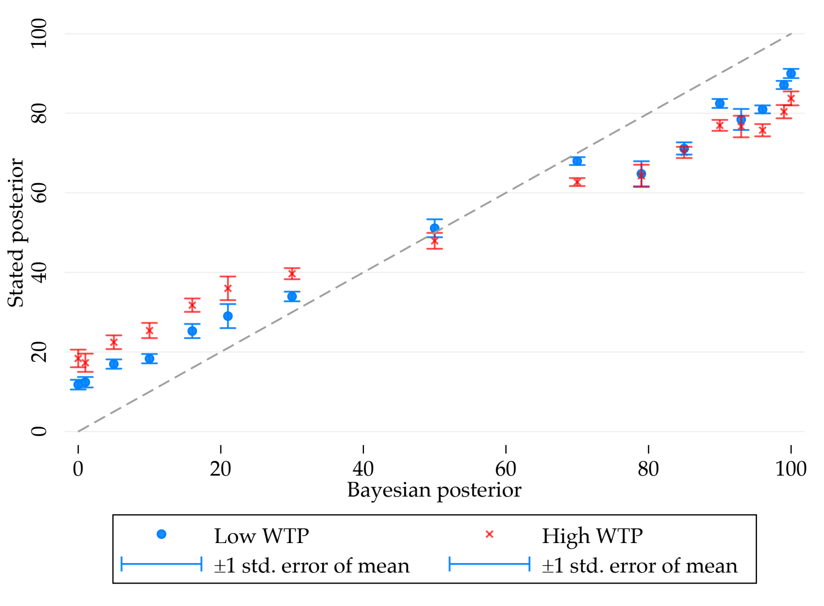

Figure 27: Relationship between stated and Bayesian posteriors, separately for subjects above / below median WTP for the Bayesian guess. The partition is done separately for each Bayesian posterior. The plot shows averages and corresponding standard error bars.

Table 11: Belief updating: Baseline tasks: WTP measure

	Dependent variable: Posterior belief			
	(1)	(2)	(3)	(4)
Bayesian posterior	0.69***	0.76***	0.76***	0.76***
	(0.01)	(0.01)	(0.01)	(0.01)
Bayesian posterior × WTP for Bayes		-0.096***	-0.096***	-0.11***
		(0.01)	(0.01)	(0.01)
WTP for Bayesian posterior		5.49***	5.47***	4.99***
		(0.76)	(0.76)	(1.01)
Session FE	No	No	Yes	No
Demographic controls	No	No	Yes	No
Subject FE	No	No	No	Yes
Observations	3187	3187	3187	3187
R^2	0.72	0.73	0.73	0.80

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.3 Grether Regressions

Table 12: Belief updating: Grether regression

	Dependent variable: Log [Posterior odds]			
Sample:	Baseline			
	(1)	(2)	(3)	(4)
Log [Likelihood ratio]	0.41***	0.44***	0.44***	0.43***
	(0.01)	(0.02)	(0.02)	(0.02)
Log [Prior odds]	0.42***	0.52***	0.52***	0.54***
	(0.02)	(0.03)	(0.03)	(0.04)
Log [Likelihood ratio] × Cognitive uncertainty		-0.16***	-0.16***	-0.15***
		(0.04)	(0.04)	(0.05)
Log [Prior odds] × Cognitive uncertainty		-0.34***	-0.35***	-0.39***
		(0.07)	(0.07)	(0.08)
Cognitive uncertainty		-0.14**	-0.16**	-0.19
		(0.07)	(0.07)	(0.13)
Session FE	No	No	Yes	No
Demographic controls	No	No	Yes	No
Subject FE	No	No	No	Yes
Observations	3104	3104	3104	3104
R^2	0.62	0.63	0.63	0.72

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes the baseline tasks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.4 Sample Size Effects

Table 13: Belief updating: Sample size effects

	Dependent variable: Cognitive uncertainty			
	(1)	(2)	(3)	(4)
Sample size	0.023***	0.023***	0.015**	0.016**
	(0.01)	(0.01)	(0.01)	(0.01)
Sample difference FE	Yes	Yes	No	No
Bayesian posterior FE	No	No	Yes	Yes
Session FE	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes
Observations	3187	3187	3187	3187
R^2	0.02	0.05	0.05	0.08

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes the baseline tasks. The sample difference fixed effects are fixed effects for the difference between the number of black and white balls that were randomly drawn. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.5 Results with Full Sample

Table 14: Belief updating: Regression analyses (full sample)

	Dependent variable:											
	Posterior belief						Log [Posterior odds]					
Sample:	Baseline		Compound		Default		Baseline		Compound		Default	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Bayesian posterior	0.75***	0.75***	0.67***	0.69***	0.54***	0.54***						
	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)						
Bayesian posterior × Cognitive uncertainty	-0.46***	-0.45***		-0.068								
	(0.05)	(0.05)		(0.06)								
Cognitive uncertainty	19.9***	19.9***		-1.12			-0.14*	-0.15*		-0.22**		-0.26***
	(3.10)	(3.12)		(3.63)			(0.08)	(0.08)		(0.10)		(0.08)
Bayesian posterior × 1 if compound problem			-0.69***	-0.68***								
			(0.04)	(0.04)								
1 if compound problem			34.5***	34.7***					-0.046	-0.013		
			(2.17)	(2.19)					(0.06)	(0.06)		
0 if *Baseline*, 1 if *Low Default*					-6.94***	-7.41***					-0.41***	-0.46***
					(0.97)	(1.02)					(0.06)	(0.06)
Log [Likelihood ratio]							0.41***	0.41***	0.40***	0.45***	0.31***	0.37***
							(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)
Log [Prior odds]							0.49***	0.49***			0.41***	0.54***
							(0.03)	(0.03)			(0.02)	(0.03)
Log [Likelihood ratio] × Cognitive uncertainty							-0.21***	-0.21***		-0.17***		-0.21***
							(0.05)	(0.05)		(0.05)		(0.04)
Log [Prior odds] × Cognitive uncertainty							-0.45***	-0.45***				-0.43***
							(0.08)	(0.08)				(0.07)
Log [Likelihood ratio] × 1 if compound problem									-0.32***	-0.30***		
									(0.03)	(0.03)		
Session FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Observations	3310	3310	2056	2056	5668	5668	3222	3222	1954	1954	5473	5473
R^2	0.59	0.60	0.45	0.46	0.44	0.44	0.50	0.50	0.40	0.41	0.42	0.44

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. In columns (1)–(2) and (7)–(8), the sample includes the baseline tasks. In columns (3)–(4) and (9)–(10), the sample includes the baseline and compound tasks (where the sample of baseline tasks is restricted to the same probabilities as in the compound tasks). In columns (5)–(6) and (11)–(12), the sample includes the low and high default tasks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.6 Results excluding Speeders

Table 15: Belief updating: Regression analyses (excl. speeders)

	Dependent variable:											
	Posterior belief						Log [Posterior odds]					
Sample:	Baseline		Compound		Default		Baseline		Compound		Default	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Bayesian posterior	0.75***	0.76***	0.68***	0.70***	0.55***	0.55***						
	(0.02)	(0.02)	(0.02)	(0.03)	(0.01)	(0.01)						
Bayesian posterior × Cognitive uncertainty	-0.46***	-0.46***		-0.094								
	(0.05)	(0.05)		(0.06)								
Cognitive uncertainty	21.4***	21.6***		1.05			-0.085	-0.085		-0.13		-0.19**
	(3.18)	(3.20)		(3.85)			(0.09)	(0.09)		(0.10)		(0.08)
Bayesian posterior × 1 if compound problem			-0.68***	-0.67***								
			(0.04)	(0.04)								
1 if compound problem			33.9***	33.9***					-0.071	-0.050		
			(2.25)	(2.30)					(0.06)	(0.06)		
0 if *Baseline*, 1 if *Low Default*					-6.64***	-6.93***					-0.40***	-0.43***
					(0.98)	(1.04)					(0.06)	(0.06)
Log [Likelihood ratio]							0.41***	0.41***	0.41***	0.45***	0.32***	0.37***
							(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)
Log [Prior odds]							0.51***	0.50***			0.43***	0.55***
							(0.04)	(0.04)			(0.02)	(0.03)
Log [Likelihood ratio] × Cognitive uncertainty							-0.22***	-0.22***		-0.19***		-0.20***
							(0.06)	(0.05)		(0.05)		(0.04)
Log [Prior odds] × Cognitive uncertainty							-0.44***	-0.44***				-0.42***
							(0.08)	(0.08)				(0.07)
Log [Likelihood ratio] × 1 if compound problem									-0.31***	-0.29***		
									(0.03)	(0.03)		
Session FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Observations	3006	3006	1874	1874	5107	5107	2925	2925	1779	1779	4930	4930
R^2	0.61	0.61	0.46	0.46	0.45	0.46	0.51	0.51	0.40	0.42	0.44	0.46

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. In columns (1)–(2) and (7)–(8), the sample includes the baseline tasks. In columns (3)–(4) and (9)–(10), the sample includes the baseline and compound tasks (where the sample of baseline tasks is restricted to the same probabilities as in the compound tasks). In columns (5)–(6) and (11)–(12), the sample includes the low and high default tasks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Additional Details and Analyses for Survey Expectations

E.1 Wording of questions

The question about the income distribution reads as:

> *Assume that in 2018, we randomly picked a household in the United States. What do you think is the probability that this household earned less than USD y in 2018, before taxes and deductions?*

Beliefs about the performance of the stock market were elicited as:

> *The S&P 500 is an American stock market index that includes 500 of the largest companies based in the United States. We randomly picked a year X between 1980 and 2018. Imagine that someone invested $100 into the S&P 500 at the beginning of year X. What do you think is the probability that, at the end of that same year, the value of the investment was less than $y? (In other words, what do you think is the probability that the S&P 500 [lost more than z% of its value / gained less than z%, or decreased in value]?*

E.2 Additional Figures

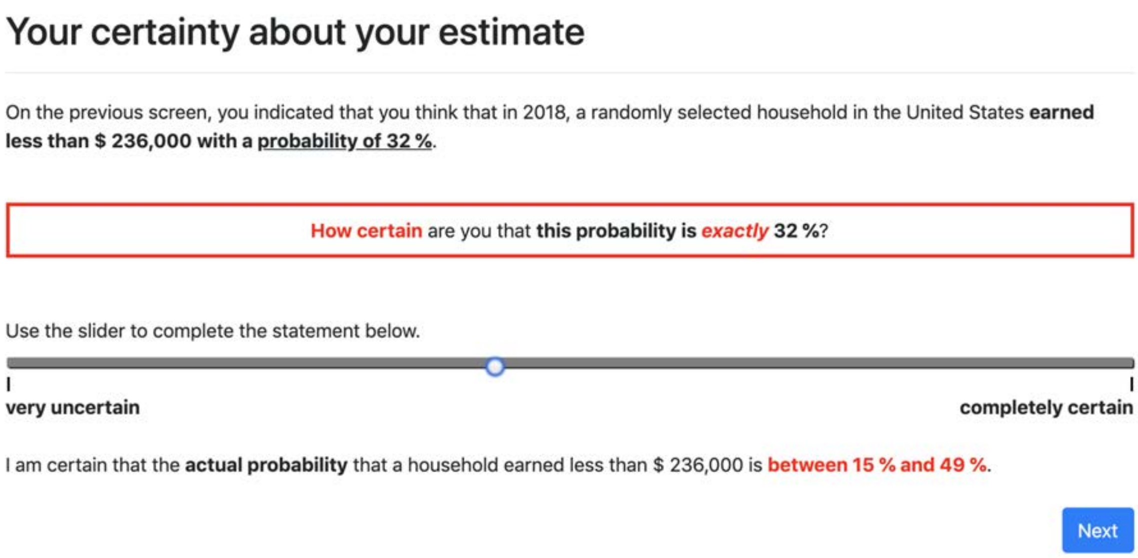

Figure 28: Decision screen to elicit cognitive uncertainty in survey expectations

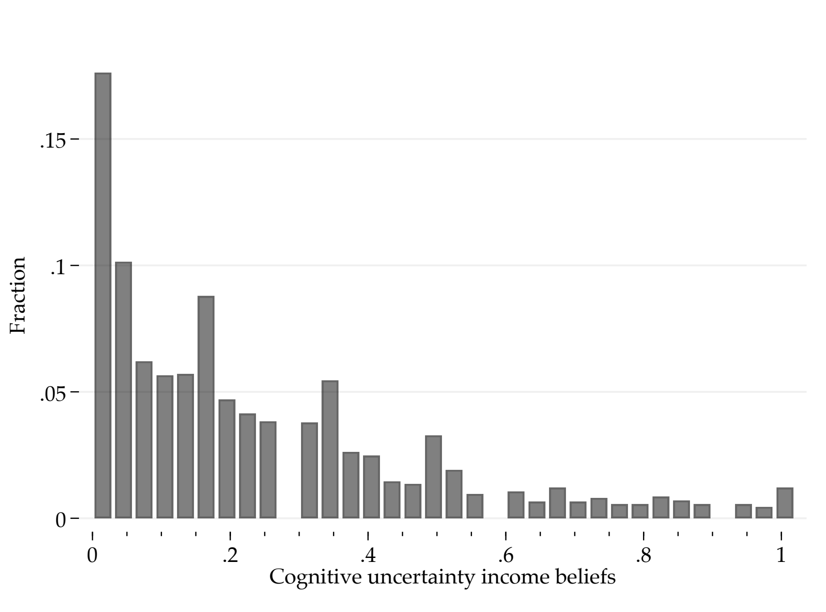

Figure 29: Histogram of cognitive uncertainty in survey expectations about income distribution

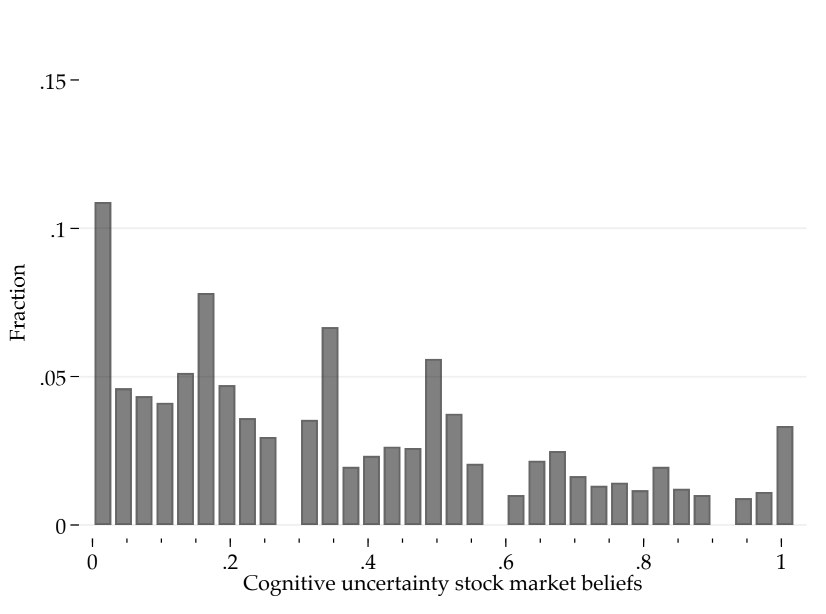

Figure 30: Histogram of cognitive uncertainty in survey expectations about the stock market

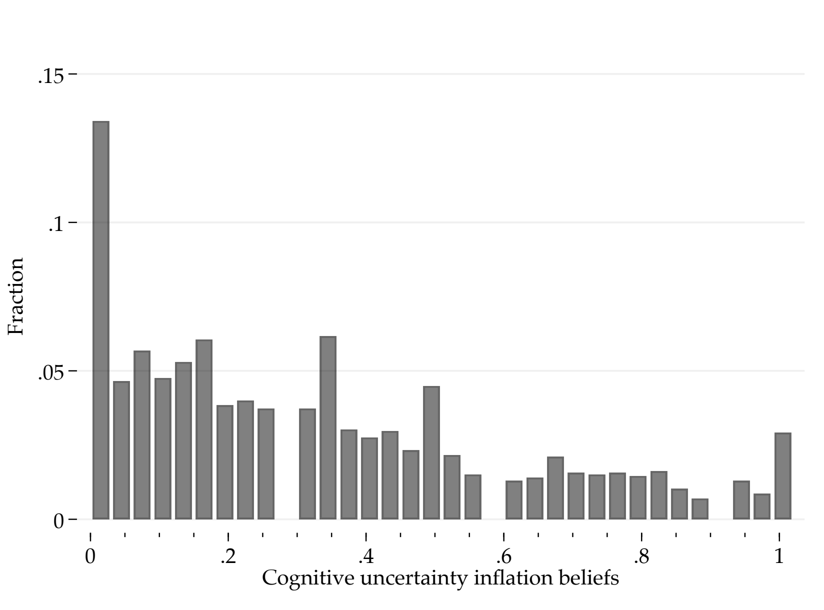

Figure 31: Histogram of cognitive uncertainty in survey expectations about inflation rates

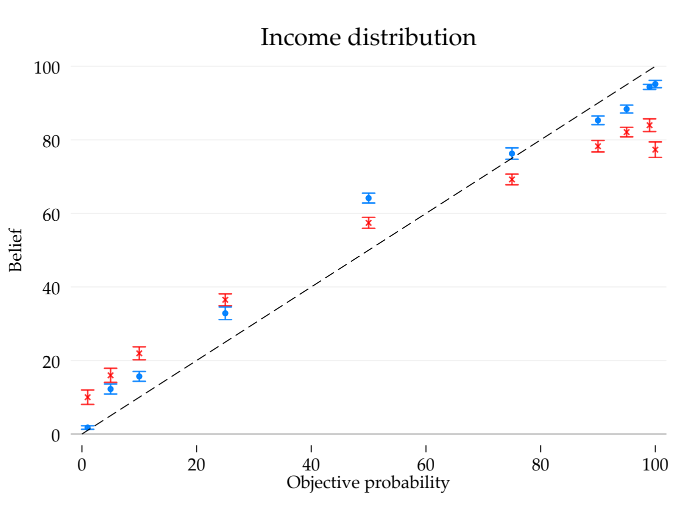

Figure 32: Survey beliefs as a function of objective probabilities, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability bucket. The question asks for the probability that a randomly selected U.S. household earns less than \$x ($N = 1,974$).

E.3 Additional Tables

Table 16: Survey expectations and cognitive uncertainty

	Dependent variable: Probability estimate about:					
	Income distr.		Stock market		Inflation rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Objective probability	0.90***	0.90***	0.69***	0.69***	0.76***	0.76***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Objective probability × Cognitive uncertainty	-0.41***	-0.41***	-0.53***	-0.52***	-0.60***	-0.60***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Cognitive uncertainty	18.9***	18.6***	24.2***	24.6***	27.5***	27.0***
	(2.37)	(2.41)	(2.27)	(2.30)	(2.86)	(2.89)
Session FE	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes
Observations	1980	1980	1892	1892	1848	1848
R^2	0.83	0.84	0.52	0.53	0.54	0.54

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. In columns (1)–(2), the question about income distribution asks participants for the probability that a randomly selected U.S. household earns less than $x. In columns (3)–(4), the question about the stock market asks participants for the probability that in a randomly selected year the S&P500 increased by less than x%. In columns (5)–(6), the question about inflation rates asks participants for the probability that in a randomly selected year the inflation rate was less than x%. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E.4 Results with Full Sample

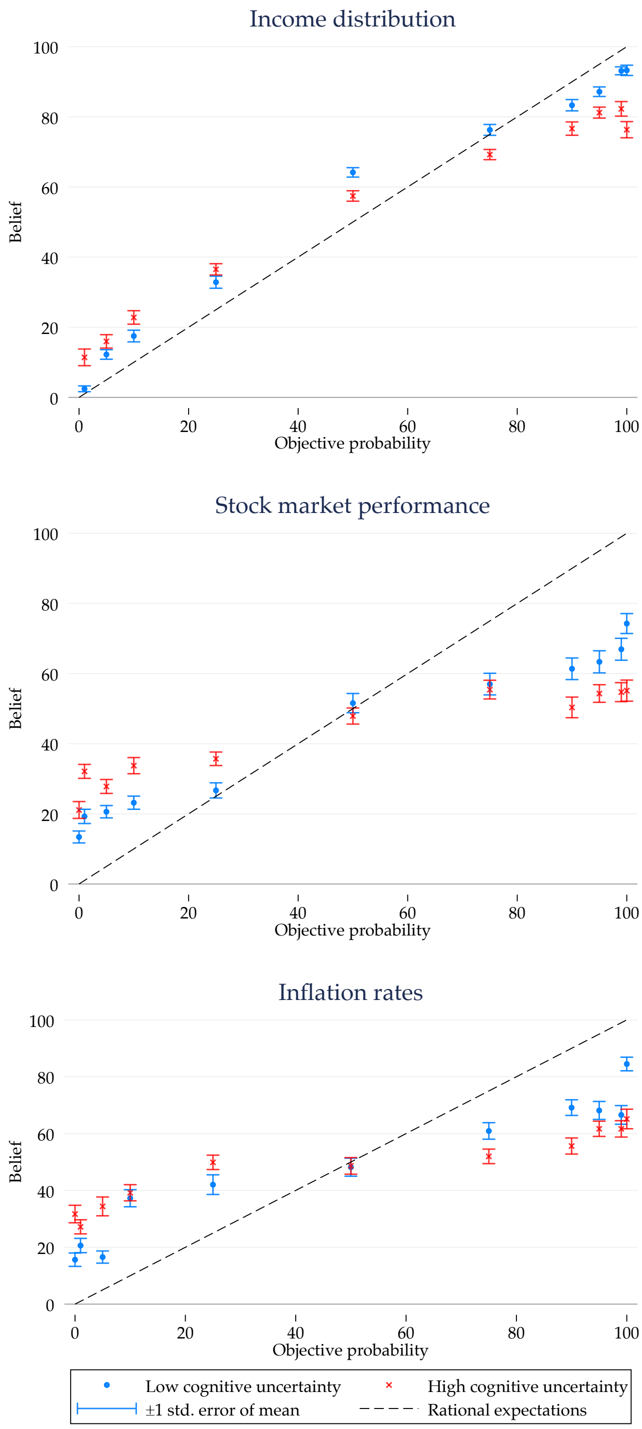

Figure 33: Survey beliefs as a function of objective probabilities, separately for subjects above / below average cognitive uncertainty (full sample). The partition is done separately for each probability bucket. In the top panel, the question asks for the probability that a randomly selected U.S. household earns less than $x. In the middle panel, the question asks for the probability that in a randomly selected year the S&P500 increased by less than x%. In the bottom panel, the question asks for the probability that in a randomly selected year the inflation rate was less than x%. $N = 2,000$ observations each.

E.5 Results excluding Speeders

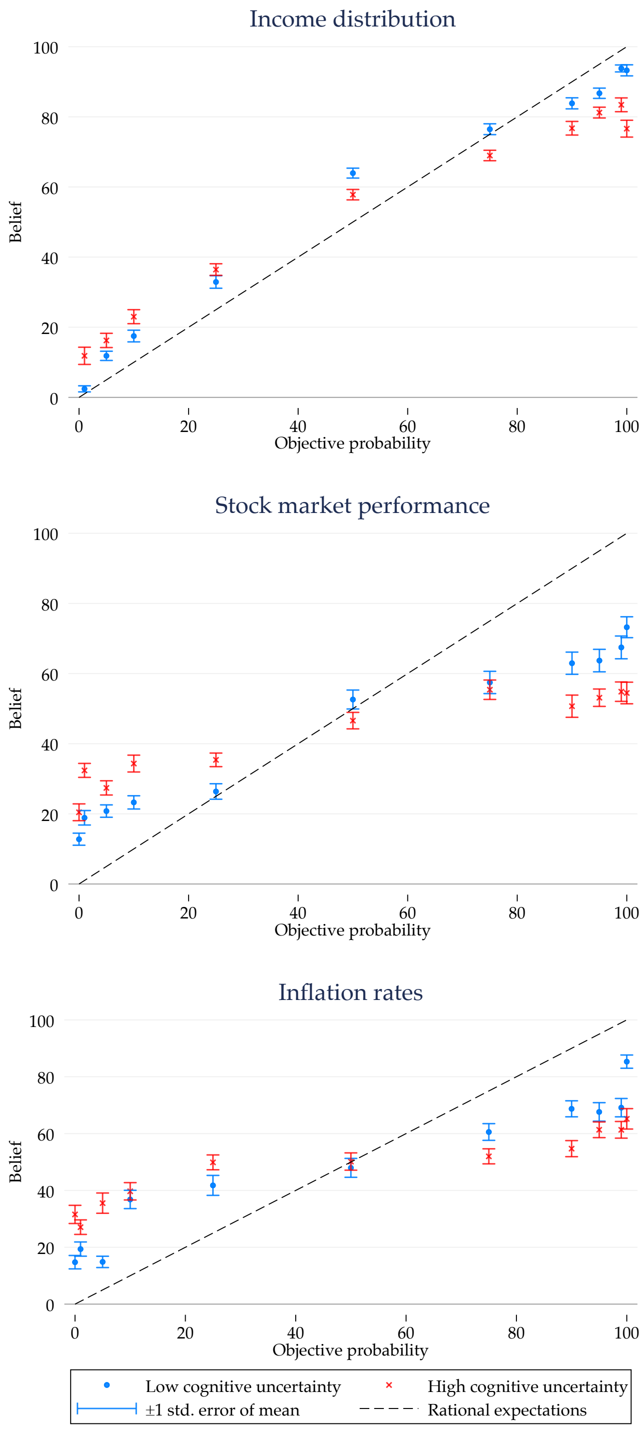

Figure 34: Survey beliefs as a function of objective probabilities, separately for subjects above / below average cognitive uncertainty (excl. speeders). The partition is done separately for each probability bucket. In the top panel, the question asks for the probability that a randomly selected U.S. household earns less than $x. In the middle panel, the question asks for the probability that in a randomly selected year the S&P500 increased by less than x%. In the bottom panel, the question asks for the probability that in a randomly selected year the inflation rate was less than x%. $N = 1,896$ observations each.

F Forward-Looking Survey Expectations

In Section 5 in the main text, we elicited respondents' survey expectations about economic variables with respect to past values, which allowed us to easily incentivize responses. In a pre-registered robustness check, we implemented the same type of survey questions, but now regarding future values of these variables. These questions are hence theoretically more appropriate in that they elicit actual expectations, but they are not financially incentivized. The sample size is $N = 400$ for each of the three domains. We apply the same criteria regarding the exclusions of outliers as in Section 5.

The results are shown in Figure 35. Here, we define "objective probabilities" based on historical data, akin to Figure 8 in the main text. The results are almost identical to those reported in the main text.

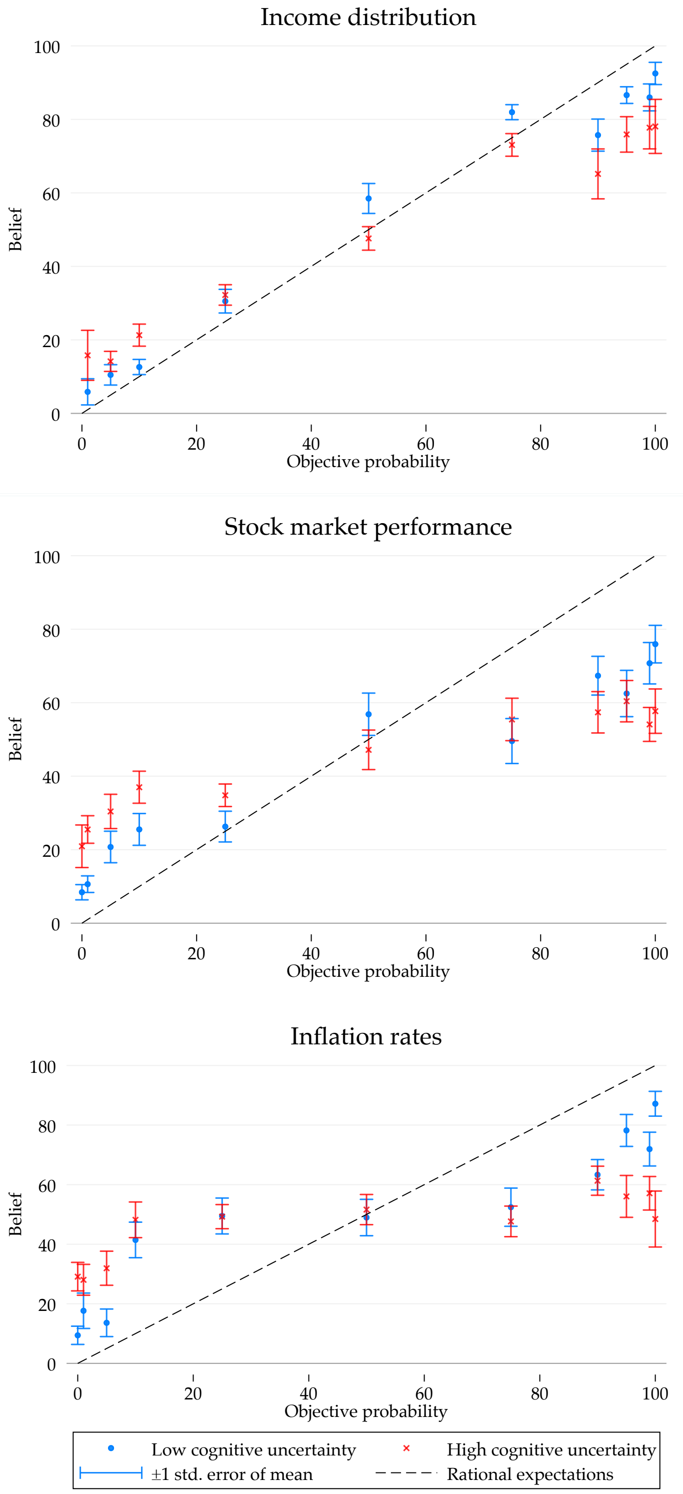

Figure 35: Survey beliefs about future variables as a function of “objective” probabilities, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability bucket. “Objective” probabilities are defined using historical data, analogously to Figure 8. In the top panel, the question asks for the probability that a randomly selected U.S. household will earn less than \$x ($N = 491$). In the middle panel, the question asks for the probability that the S&P500 will increase by less than x% ($N = 463$) over the course of one year. In the bottom panel, the question asks for the probability that the inflation rate will be than x% ($N = 478$).

G Derivation of and Estimates for Gonzalez-Wu Weighting Function

G.1 Additional Derivations for Weighting Function with Log Coding

We assume that a probability p is transformed into a quantity q in log odds space by applying

$$q = Q(p) = \ln\frac{p}{1-p}. \tag{21}$$

This means we now assume that the decision-relevant quantity is a probability in log odds space q about which an agent receives a signal $s = q+\varepsilon$. This will result in shrinkage of probabilities in log odds space:

$$q(s) = \lambda s + (1-\lambda)q^d. \tag{22}$$

In the following, we will focus on medians, which have the attractive property that for any strictly monotone function Y, $\text{Median}(Y(x)) = Y(\text{Median}(x))$. Over many draws of s, the median posterior q^e about probability p after encoding in log odds space and shrinkage is:

$$q^e(q) := \text{ Median}\left(q(s)|q\right) = \lambda q + (1-\lambda)q^d. \tag{23}$$

From this we can derive the implied median posterior probability p by applying the inverse log odds function $P(q) = Q^{-1}(q) = \frac{1}{1+e^{-q}}$:

$$p^e(p) = P(q^e) = \frac{1}{1+\exp\left(-\lambda\ln\frac{p}{1-p} - (1-\lambda)\ln\frac{p^d}{1-p^d}\right)}. \tag{24}$$

G.2 Estimation of Gonzalez-Wu function

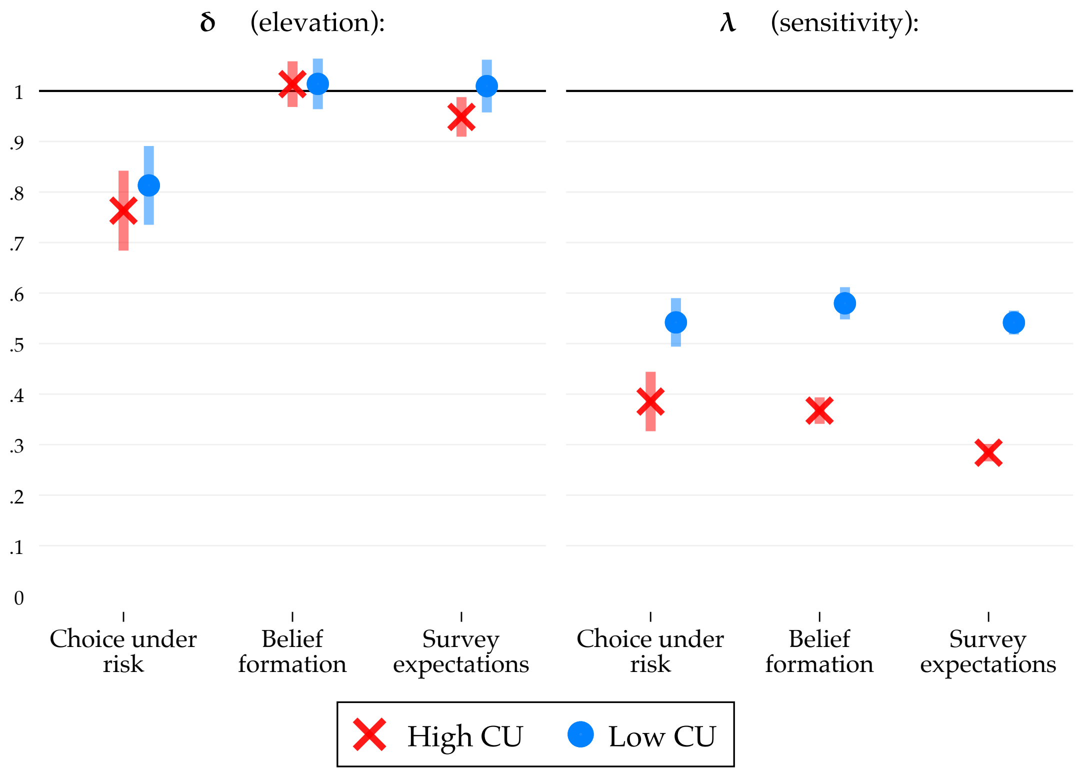

Figure 36: Estimates of equation (13), confidence bands indicate 95% confidence intervals (standard errors clustered at subject level). Cognitive uncertainty is split at sample average. The samples include the same observations as in the baseline analyses in Sections 3–5.

H Results on Stake Size Increase

H.1 Stake Size and Choice Under Risk

To manipulate the size of financial incentives, we implement a within-subjects manipulation. We implemented the same procedures as described in Section 3, except that we only used gain lotteries. Subjects completed six choice lists, one of which determined a subject's payment in case the choice under risk part of the experiment got selected for payment (probability 1/3). Across the six choice lists, the probability of being payout-relevant varied in a transparent way. On top of the decision screen, we informed subjects about the probability that this choice list would determine their payout. For five tasks, this probability was given by 1% and for one task by 95%. As a measure of cognitive effort, we recorded subjects' (log) response times. 150 subjects participated in this treatment, which was also pre-registered.

The results are reported in Table 17.[25] Exploiting variation within subjects across tasks, we find that response times increase significantly from 25 seconds on average to 36 seconds on average in the high stakes task. However, this increase in response times does not translate into a significant change in cognitive uncertainty (columns 3 and 4), nor into less compression (columns 5 and 6).

Table 17: Effects of stake size increase in choice under risk

	Dependent variable:					
	Log [Response time]		Cognitive uncertainty		Normalized CE	
	(1)	(2)	(3)	(4)	(5)	(6)
1 if high stakes	0.26***	0.26***	0.0061	0.0065	0.41	-1.65
	(0.06)	(0.06)	(0.02)	(0.02)	(2.27)	(2.16)
Probability of payout					0.69***	0.68***
					(0.03)	(0.03)
Probability of payout × 1 if high stakes					0.022	0.065*
					(0.04)	(0.04)
Subject FE	No	Yes	No	Yes	No	Yes
Observations	893	893	893	893	893	893
R^2	0.02	0.50	0.00	0.52	0.60	0.79

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[25] We again apply the same outlier exclusion criteria as in the main text.

H.2 Stake Size and Belief Updating

To manipulate the size of financial incentives, we again implement a within-subjects manipulation. We implemented the same procedures as described in Section 4, except that we did not elicit the WTP for the optimal guess. Subjects completed six updating tasks, one of which determined a subject's payment in case the belief updating part of the experiment got selected for payment (probability 1/3). Across the six tasks, the probability of being payout-relevant varied in a transparent way. On top of the decision screen, we informed subjects about the probability that this task would determine their payout. For five tasks, this probability was given by 1% and for one task by 95%. As a measure of cognitive effort, we recorded subjects' (log) response times. 150 subjects participated in this treatment, which was also pre-registered.

The results are reported in Table 18.[26] Exploiting variation within subjects across tasks, we find that response times increase significantly. Cognitive uncertainty decreases (columns 3 and 4) and people respond more to the Bayesian posterior (columns 5 and 6), but only very mildly so.

Table 18: Effects of stake size increase in belief updating

	Dependent variable:					
	Log [Response time]		Cognitive uncertainty		Posterior belief	
	(1)	(2)	(3)	(4)	(5)	(6)
1 if high stakes	0.19***	0.19***	-0.025*	-0.026	-2.76	-3.56
	(0.06)	(0.07)	(0.01)	(0.02)	(2.53)	(2.64)
Bayesian posterior					0.59***	0.58***
					(0.03)	(0.03)
Bayesian posterior × 1 if high stakes					0.065	0.080*
					(0.04)	(0.04)
Subject FE	No	Yes	No	Yes	No	Yes
Observations	869	869	869	869	869	869
R^2	0.01	0.46	0.00	0.50	0.61	0.70

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[26] We again apply the same outlier exclusion criteria as in the main text.

I Censoring and Random Choice

I.1 General Discussion

Below we discuss, first, evidence on censoring and, second, the role of random choice in generating our findings.

In all our experiments, the use of a bounded response scale can lead to censoring of both the choice or belief that a subject states and the range of cognitive uncertainty indicated using the slider. This may affect the observed relationship between actions and cognitive uncertainty in two ways. First, choices and beliefs may be influenced by boundary effects. Assume, for example, that a subject in a belief updating task wants to state a true posterior belief of 95%. However, some form of decision noise such as trembling when submitting a response leads her to instead indicate a posterior belief that is uniformly drawn from within ±10% of her true posterior, i.e., she would end up with any posterior between 85% and 105% with equal probability. Since it is not possible to state a posterior greater than 100%, she will state 100% whenever she would like to state something greater than 100%, leading to an observed posterior that is *lower* than 95% in expectation. Importantly, this distortion in observed beliefs away from the boundary will be greater for someone with greater decision noise. If subjects' cognitive uncertainty statements then accurately reflect the amount of trembling, i.e., the length of the trembling interval in this case, this form of censoring will mechanically generate a positive relationship between the extent of cognitive uncertainty and shrinkage. We find that the actual amount of bunching at the upper and lower bounds of the response scales, however, is small: it is 4.28% of observations in choice under risk, 2.6% of observations in belief updating, and 6.61% in survey expectations. In Appendix I.2 below, we show that the observed relationship between cognitive uncertainty and choices or beliefs is virtually unaffected when excluding these observations. Moreover, note that we generally observe a pronounced relationship between cognitive uncertainty and actions for probabilities far away from the boundaries, e.g. at 25% and 75%. An even more generous interpretation of boundary effects is that people "shy away" from stating extreme answers, even if these would not *exactly* correspond to one of the boundaries. Decision noise would need to be (implausibly) extreme to rationalize our findings towards the middle of the probability range. In addition, recall that the random choice account is incompatible with the default manipulations reported in choice under risk and belief formation.

Second, censoring might occur when choosing an interval on the response scale to indicate cognitive uncertainty. While the interval increases symmetrically when moving the slider to the left, it increases asymmetrically once it hits one of the response scale boundaries. One may think that subjects stop moving the slider to the left once they hit a boundary. This implies that measured cognitive uncertainty tends to be smaller for responses that are closer to a boundary, again leading to a mechanical relationship between observed cognitive uncertainty and the amount of shrinkage. In our data, we find that 23.25% of cognitive uncertainty intervals in choice under risk, 25.93% in belief updating and 16.66% in survey expectations are censored at one of the boundaries. However, as we exclude those observations, the relationship between cognitive uncertainty and choices or beliefs persists in the same way as before (see Appendix I.3 below), showing that our findings are not an artifact of censoring due to bounded response scales.

Finally, note that our model of Bayesian shrinkage differs from models of random choice. Specifically, a random choice model implies that people's actions are subject to random noise, but there is no directional shrinkage in response. In that case, cognitive uncertainty might reflect the degree of decision noise. Pure decision choice in conjunction with boundary effects would create a statistical reversion of average decisions to the middle of the response scale. As discussed before, this explanation is consistent with some of our evidence, but cannot explain evidence from our default treatments. In the following, we exploit that the empirical distribution of choices (for a given payoff probability) in choice under risk sheds light on the validity of a random choice explanation as opposed to our shrinkage model. A random choice explanation of cognitive uncertainty implies that differences in observed mean actions are driven by the *tails* of the distribution. Specifically, with random choice (1) the mode of the empirical distributions of choices should be identical, (2) but because higher cognitive uncertainty (due to higher decision noise) is associated with a flatter empirical distribution, there will be more bunching close to the boundaries, which drives the effect on mean choices. By contrast, the shrinkage model predicts a shift of the entire distribution due to cognitive uncertainty, i.e., even the modal action should move in the direction of the default. Figure 37 displays histograms of normalized certainty equivalents from choice under risk, separately for each payout probability and split by above vs. below median reported cognitive uncertainty. The empirical distributions are clearly in line with the predictions of our shrinkage model, but at odds with pure random choice. Specifically, note that (1) in all but one case, the modal choice under high cognitive uncertainty differs (in the predicted direction of 50) from that under low cognitive uncertainty, and (2) for each payout probability, there is *less* bunching at the boundaries among observations with higher cognitive uncertainty, rather than more bunching as implied by random choice.

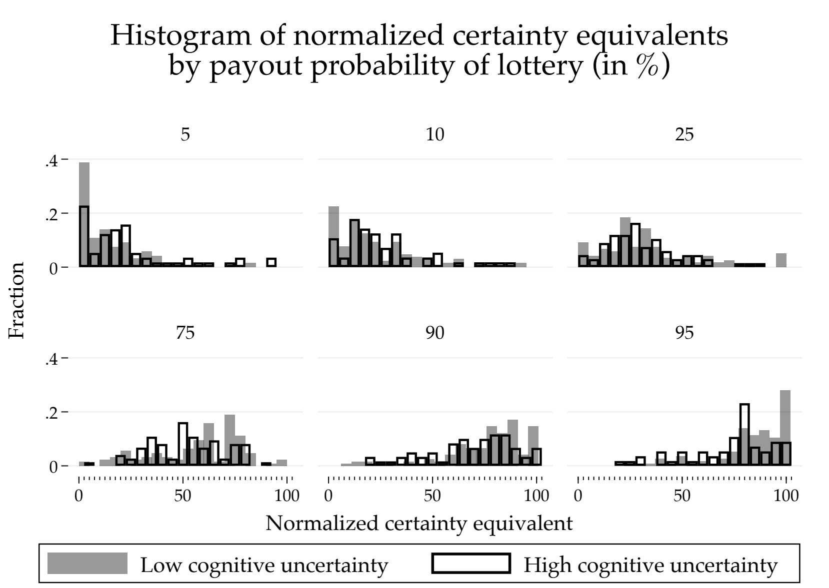

Figure 37: Histogram of normalized certainty equivalents for each payout probability, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability bucket. Normalized certainty equivalents (implied probability weights) are computed as certainty equivalent divided by payout probability. The figure is based on 1,102 certainty equivalents.

I.2 Censored Choices and Beliefs

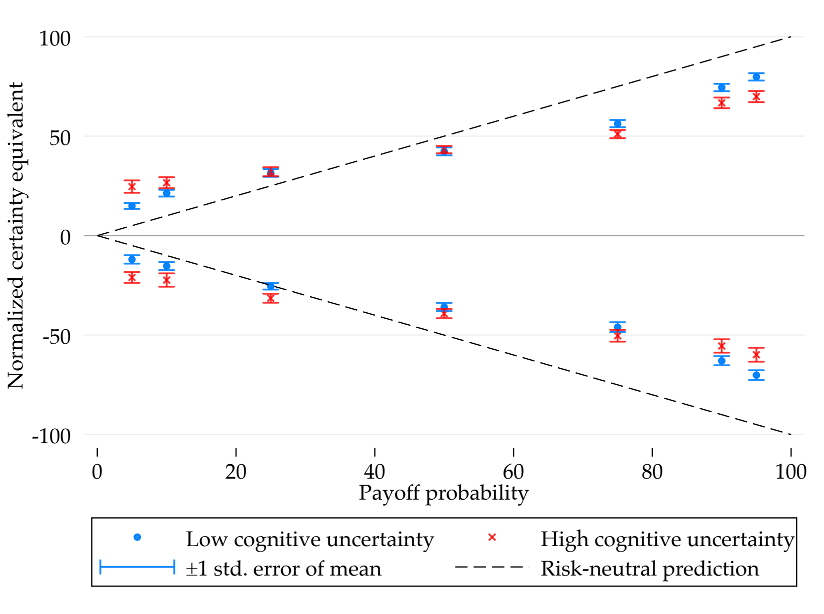

Figure 38: Probability weighting function excluding censored choices, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability × gains / losses bucket. The plot shows averages and corresponding standard error bars. Normalized certainty equivalents (implied probability weights) are computed as certainty equivalent divided by payout probability. The figure excludes 4.28% of the original data that is based on 2,525 certainty equivalents of 700 subjects.

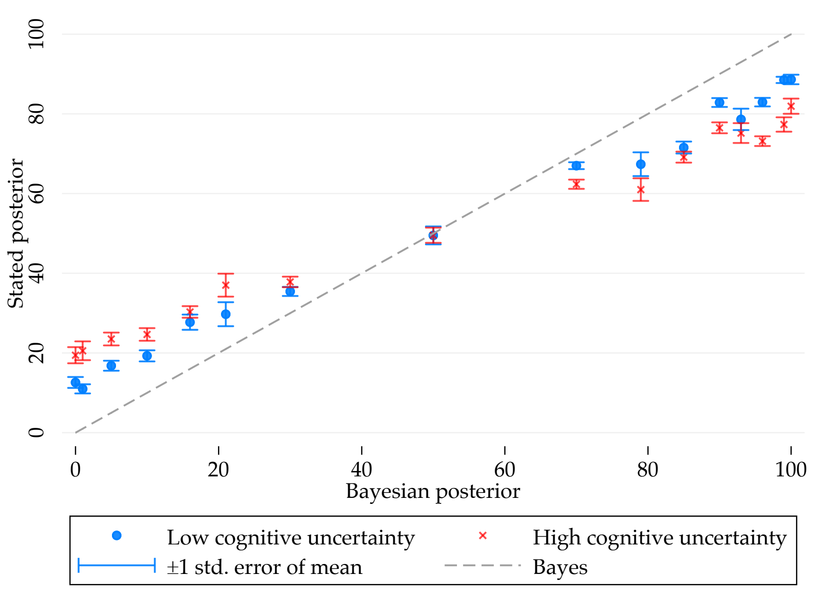

Figure 39: Relationship between average stated and Bayesian posteriors after excluding censored beliefs, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each Bayesian posterior. Bayesian posteriors are rounded to the nearest integer. We only show buckets with at least 15 observations. The figure excludes 2.6% of the original data that is based on 3,187 beliefs of 700 subjects.

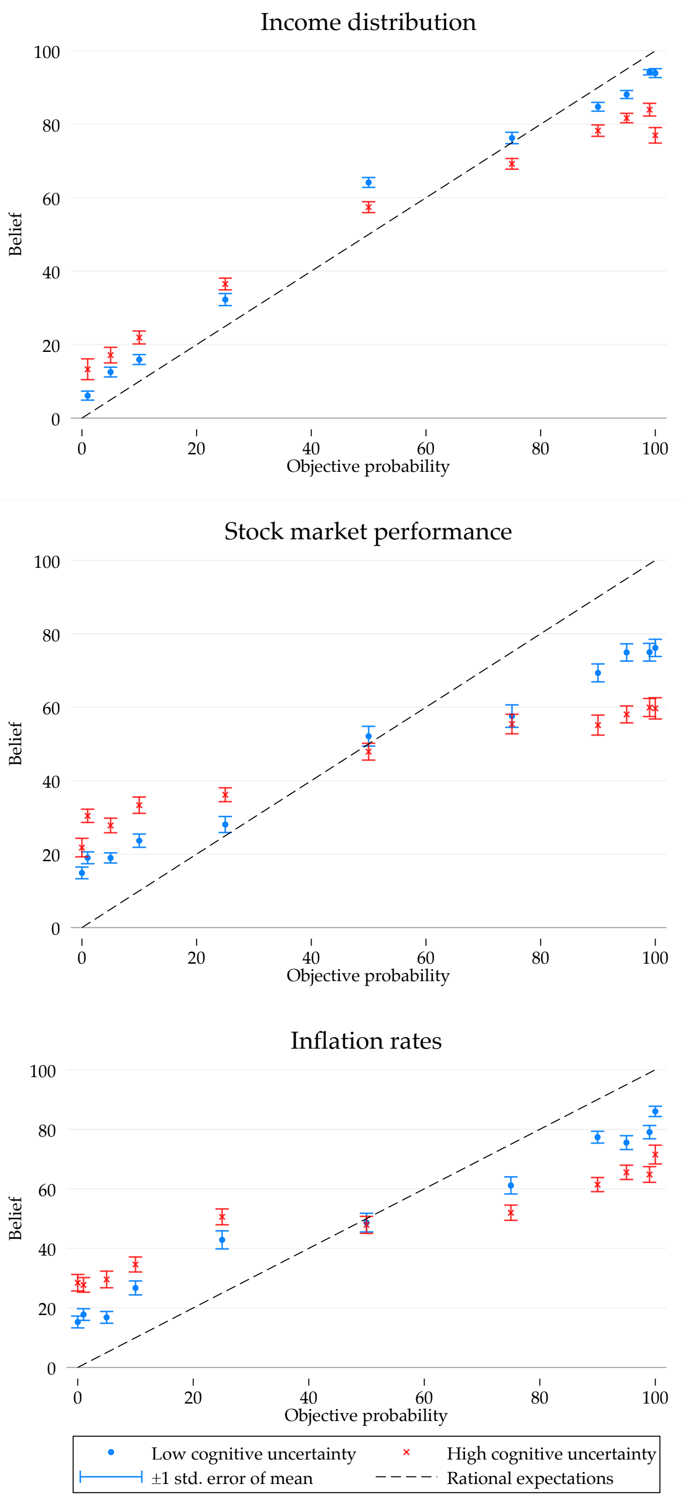

Figure 40: Survey beliefs as a function of objective probabilities after excluding censored beliefs, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability bucket. In the top panel, the question asks for the probability that a randomly selected U.S. household earns less than $x. In the middle panel, the question asks for the probability that in a randomly selected year the S&P500 increased by less than x%. In the bottom panel, the question asks for the probability that in a randomly selected year the inflation rate was less than x%. The figure excludes 6.61% of the original data that is based on 5,703 observations.

I.3 Censored Cognitive Uncertainty Range

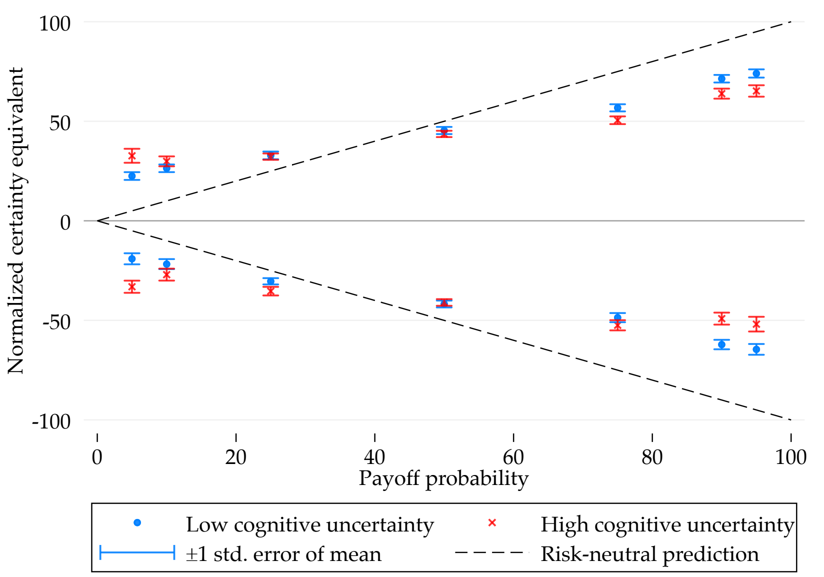

Figure 41: Probability weighting function excluding censored cognitive uncertainty ranges, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability × gains / losses bucket. The plot shows averages and corresponding standard error bars. Normalized certainty equivalents (implied probability weights) are computed as certainty equivalent divided by payout probability. The figure excludes 23.25% of the original data that is based on 2,525 certainty equivalents from 700 subjects.

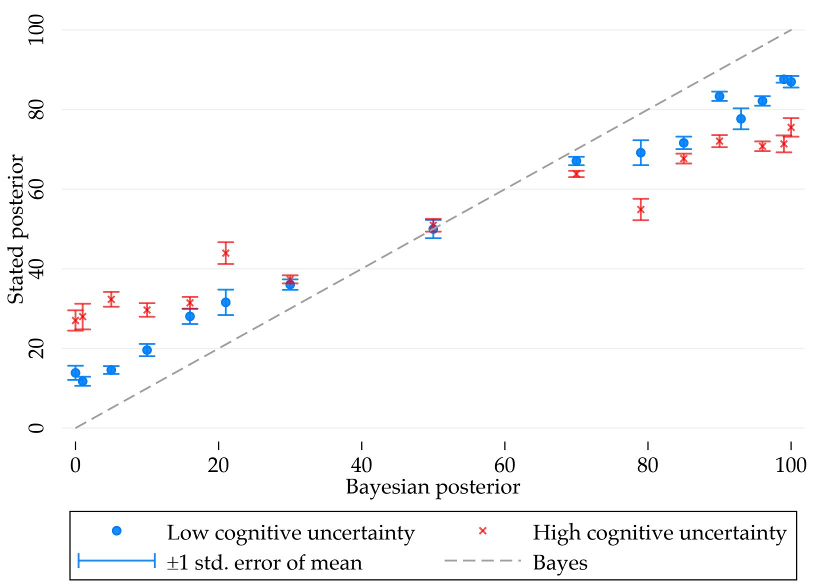

Figure 42: Relationship between average stated and Bayesian posteriors after excluding censored cognitive uncertainty ranges, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each Bayesian posterior. Bayesian posteriors are rounded to the nearest integer. We only show buckets with at least 15 observations. The figure excludes 25.93% of the original data that is based on 3,187 beliefs of 700 subjects.

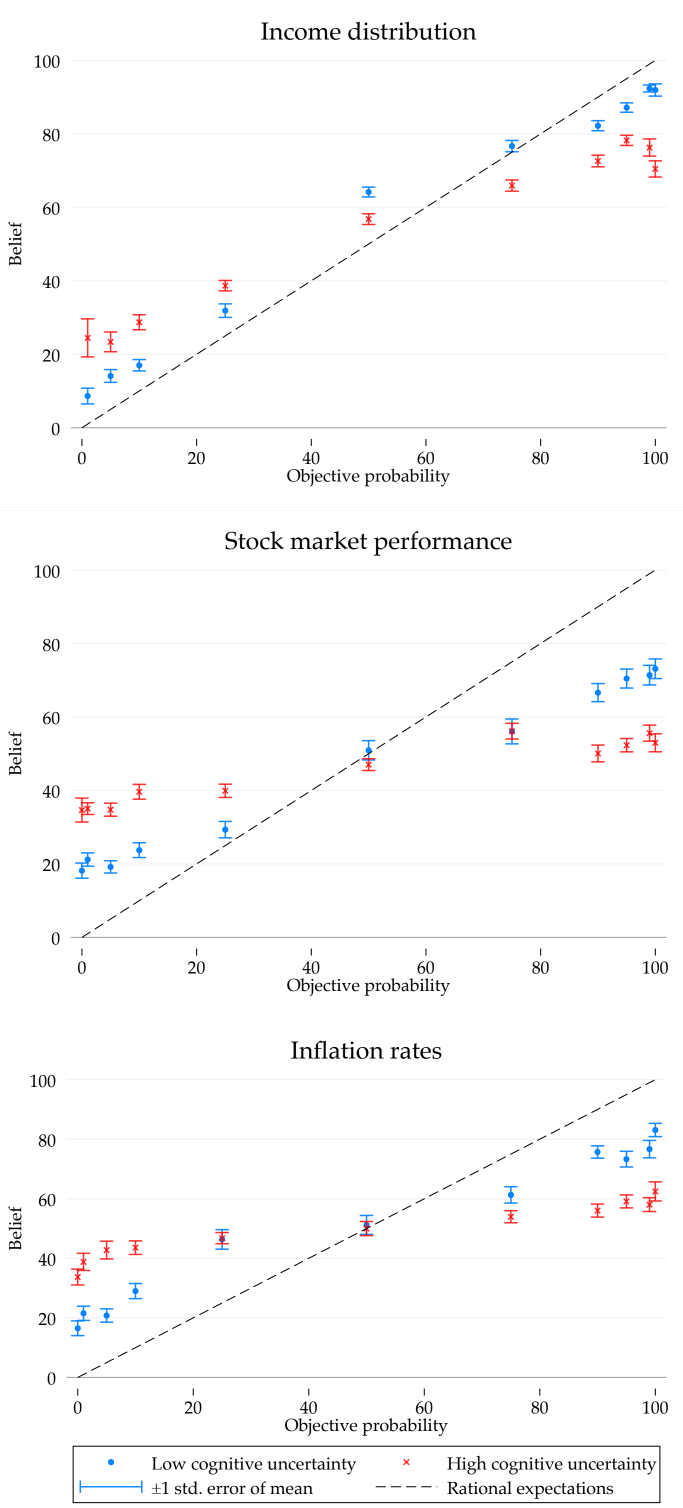

Figure 43: Survey beliefs as a function of objective probabilities after excluding censored cognitive uncertainty ranges, separately for subjects above / below average cognitive uncertainty. The partition is done separately for each probability bucket. In the top panel, the question asks for the probability that a randomly selected U.S. household earns less than $x. In the middle panel, the question asks for the probability that in a randomly selected year the S&P500 increased by less than x%. In the bottom panel, the question asks for the probability that in a randomly selected year the inflation rate was less than x%. The figure excludes 16.66% of the original data that is based on 5,703 observations.

J Additional Ambiguity Experiment

In addition to the experiments reported in Section 3, we implemented an additional set of pre-registered ambiguity experiments. These experiments delivered statistically significant results in line with our pre-registered predictions. However, as explained below, we now believe that these experiments are conceptually less-than-ideal from the perspective of our framework, which is why we relegate them to an Appendix.

J.1 Experimental Design

The basic design builds on Dimmock et al. (2015) and aims at eliciting matching probabilities for ambiguous lotteries. In a given choice list, the left-hand side option A was constant and given by an ambiguous lottery. The ambiguous lottery was described as random draw from an urn that comprises 100 balls of ten different colors, where the precise composition of colors is unknown. A known number of these colors n were "winning colors" that resulted in the same payout \$x, while other colors resulted in a zero payout. Option B, on the right-hand side, varied across rows in the choice list and was also given by a lottery with upside \$x. Here, the number of "winning balls" was known and varied from 0% to 99% in 3% steps. Subjects were always given the option to pick their preferred winning colors.

A subject completed six choice lists, where the payout $x \in \{15, 20, 25\}$ and the number of winning colors $n \in \{1, 2, \ldots, 9\}$ were randomly determined. Before each decision screen, subjects were always given the opportunity to pick their winning colors.

Cognitive uncertainty was measured analogously to choice under risk. After subjects had indicated their probability equivalent range for an ambiguous lottery, the subsequent screen asked them how certain they are that this range actually corresponds to how much the lottery is worth to them. Operationally, subjects used a slider to calibrate the statement "I am certain that to me the lottery is worth as much as playing a lottery over \$x with a known number of between x and y winning balls." 200 AMT workers participated in these experiments and earned an average of \$7.20.

J.2 Results

In the baseline analysis, we again exclude extreme outliers, defined as matching probability strictly larger than 75% for at most two winning colors, and matching probability strictly smaller than 25% for more than eight winning colors. This is the case for 1.6% of our data. We find that the response function of subjects with higher cognitive uncertainty is significantly less sensitive to variation in the number of winning colors (shallower), see the regressions in Table 19. This reduction in sensitivity corresponds to our

Table 19: Insensitivity to ambiguous "likelihood" and cognitive uncertainty

	Dependent variable: Matching probability	
	(1)	(2)
Number of winning colors * 10	0.63***	0.63***
	(0.04)	(0.04)
Number of winning colors * 10 × Cognitive uncertainty	-0.12	-0.11
	(0.11)	(0.12)
Cognitive uncertainty	1.39	-0.35
	(4.76)	(4.75)
Session FE	No	Yes
Demographic controls	No	Yes
Observations	1181	1181
R^2	0.49	0.51

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is a subject's matching probability, computed as midpoint of the switching interval. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

main hypothesis, which is also what we re-registered. At the same time, we do not find that high cognitive uncertainty subjects are more ambiguity seeking than low cognitive uncertainty subjects for unlikely events.

J.3 Interpretive Problems

The analysis above focuses on whether reported matching probabilities of subjects with higher cognitive uncertainty are less sensitive to the variation in winning colors. However, our framework in Section 2 *only* makes this prediction if one assumes that the state space is binary (win-lose), so that subjects are hypothesized to "shrink" ambigious probabilities towards 50:50. However, in the experiments, the state space was represented through ten different colors, some of which are winning and some of which are losing colors. As discussed in Section 3.4, a plausible alternative view is that in this situation there are actually ten states of the world, one for each color. In this case, our framework does not predict that subjects shrink their matching probabilities towards 50:50. To see this, take the example that there are three winning colors. In this case, the ignorance prior (for winning) would be given by 30%. In other words, subjects would be hypothesized to shrink an ambiguous probability of three winning colors towards a cognitive default of 30%, which does not produce any shrinking theoretically. For this reason, we view these experiments as imperfect.

K Additional Cognitive Load Experiment

In this experiment, we manipulated cognitive load between subjects. Half of the subjects were randomly selected to work on a separate task next to filling out the price lists. Specifically, they had to sum up numbers that were flashed for 0.5 seconds in random intervals on top of the screen. Specifically, integers between 3 and 8 appeared in intervals of 1 to 5 seconds. On an additional screen that directly followed the price list, subjects had to enter their guess of the sum of the flashed numbers. We hypothesized that this task takes some attention away from the price list and increases subjects' cognitive uncertainty about their choices. To provide a financial incentive for the summation task, subjects were paid $1 if they entered the correct sum in a randomly selected round.

These additional experiments were not pre-registered. We collected a sample of $N = 169$ subjects (86 in load condition, 83 in no-load condition). The load manipulation increased average cognitive uncertainty by 15%. In line with our hypothesis, cognitive load is associated with more risk averse choices for high-probability gains and low-probability losses, but with more risk seeking choices for low-probability gains and high-probability gains. The data are displayed in Figure 44. In a regression of the normalized absolute certainty equivalent (pooling gain and loss lotteries) on the payoff probability, an indicator for the load condition, and the interaction of both, we find that the interaction term is negative and significant at the 5% level.

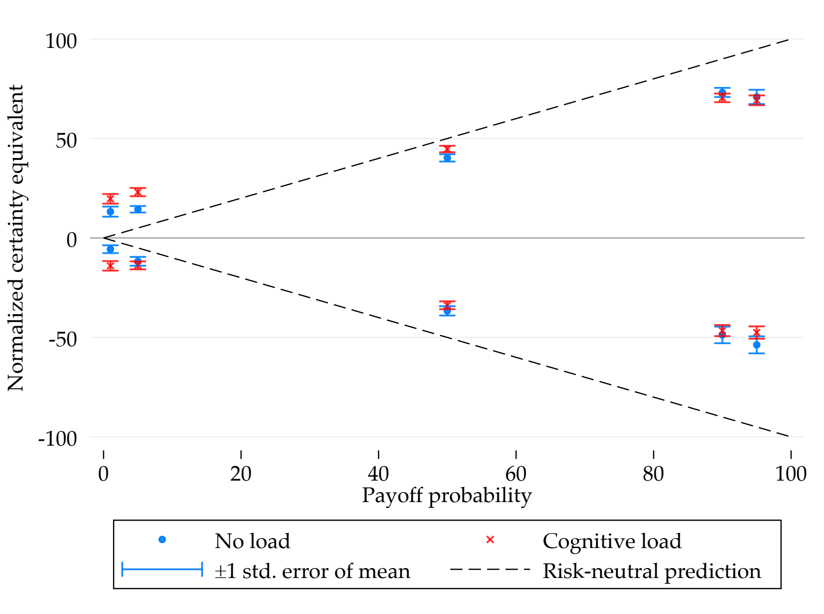

Figure 44: Probability weighting function, separately for subjects in the load condition and the no-load condition. The plot shows averages and corresponding standard error bars. Normalized certainty equivalents (implied probability weights) are computed as certainty equivalent divided by payout probability. The figure is based on $N = 86$ subjects in the load condition and $N = 83$ subjects in no-load.

L Experimental Instructions and Control Questions

L.1 Treatment *Baseline Risk*

Welcome

Thank you for participating in this study.

This study will take approximately 25 minutes to complete.

You will earn a fixed reward of $0.50. In addition, you can earn a bonus for completing the study. To complete the study, you will need to read all instructions carefully and answer the corresponding comprehension questions correctly. You receive a minimum bonus of $1.20 when you complete the study.

You can earn an ***additional* bonus**. At the end of the study, one of the tasks will be randomly selected and your decision in this task determines your additional bonus. The chance that you get paid an additional bonus for the first part is 1 in 3.

Important information

- You should think about each task **independently** of all other tasks in this study. There is no point in strategizing across tasks.
- You will note that we sometimes ask you to work on similar-sounding tasks. Theses tasks might have similar answers, or very different ones. Please consider each individual task **carefully**.

Next

Part 1: Instructions

Please read these instructions carefully. There will be comprehension checks. If you fail these, you will be excluded from the study and only receive the reward of $0.50.

In this study, there are various lotteries, all of which pay different amounts of money with different probabilities. An example lottery is:

With probability **70%: Get $ 20**
With probability **30%: Get $ 5**

This means that the lottery pays **either** $20 **or** $5 (with different probabilities), but a lottery always only pays out one of the dollar amounts. The lotteries will actually be played out by the computer and determine your earnings in exactly the way we describe on the following pages.

In total, we have created 6 lotteries. For each lottery, you will make decisions on two consecutive decision screens.

- Decision screen 1: You will make various choices to indicate **how much** a lottery is worth to you.
- Decision screen 2: You will indicate **how certain** you are about how much **exactly** a lottery is worth to you.

Throughout the experiment, there are no right or wrong answers, because how much you like a lottery depends on your personal taste. We are only interested in learning about what you prefer.

Decision screen 1

On decision screen 1, you will be asked to choose which of two payment options you prefer. You will see choice lists such as the one below, where each row is a separate choice. In every list, the left-hand side option (Option A) is one of the lotteries that we generated, which is identical in all rows. The right-hand side option (Option B) is a safe payment that you would receive *with certainty*, so there is no risk attached. This safe payment in Option B increases as you go down the list.

You should consider the choice in each row **independently** of all other rows, because in the end one row will be randomly selected for payout and you will receive the option that you selected in that row. In some choice lists that you complete, you will receive a budget. If so, this budget will be listed at the top of your screen and paid out to you along with the payouts from the lottery.

Usually, people start by preferring Option A for small certain amounts (at the top of the list). At some point they switch to Option B as they proceed down the list, because the certain amount associated with Option B increases, so that Option B becomes more attractive. **Thus, an effective way to complete these choice lists is to determine in which row you would like to switch from Option A to Option B.**

Based on your decisions in this choice list, we assess **how much the lottery is worth to you by using those certain payments where you switch from Option A to Option B.** For example, in the example choice list below, the lottery would be worth between $13 and $14 to you, because this is where switching occurs.

Option A		Option B
With probability **70%**: Get $ 20 With probability **30%**: Get $ 5		With certainty: Get $ 0
		With certainty: Get $ 1
		With certainty: Get $ 2
		With certainty: Get $ 3
		With certainty: Get $ 4
		With certainty: Get $ 5
		With certainty: Get $ 6
		With certainty: Get $ 7
		With certainty: Get $ 8
		With certainty: Get $ 9
		With certainty: Get $ 10
		With certainty: Get $ 11
		With certainty: Get $ 12
		With certainty: Get $ 13
		With certainty: Get $ 14
		With certainty: Get $ 15
		With certainty: Get $ 16
		With certainty: Get $ 17
		With certainty: Get $ 18
		With certainty: Get $ 19
		With certainty: Get $ 20

Click "Next" to read about decision screen 2.

Next

Decision screen 2 (for the same lottery)

For any given lottery, you may actually be **uncertain about how much money it is really worth** to you. For some lotteries, you may know exactly how much they are worth to you. For other lotteries, you may feel uncertain about whether a lottery is worth, say, $12, $13, $14, $15, or $16 to you.

On the second decision screen for a given lottery, we will hence ask you to use a slider to indicate **how certain** you are that your decisions on the first decision screen correspond **exactly** to how much the lottery is **worth to you**.

- The slider is linked to the statement below the slider. As you move around the slider, the values in this statement change. Your task is to position the slider in such a way that the statement corresponds to how you feel about the lottery.
- The **further you move the slider to the right, the more certain** you are about how much exactly the lottery is worth to you.
- For every step that you move the slider **further to the left (the less certain you are)**, the range of dollar values that you consider possible increases by $0.50.

Example

Suppose that on the first decision screen you indicated that the lottery below is worth between $13 and $14 to you (because you switched from Option A to B between $13 and $14). Your screen would then look like this:

Example lottery:

With probability **70%: Get $ 20**
With probability **30%: Get $ 5**

How certain are you that to you this lottery **is worth *exactly* the same as getting between $13 and $14** for sure?

very uncertain **completely certain**

I am certain that the lottery is worth **Use the slider!** to me.

Next

Your payment for part 1

If this part is randomly selected for payment in the end, your *additional* bonus will be determined as follows:

The computer will randomly select one choice list for payout, with equal probability. We then randomly select one of your decisions from this choice list (with equal probability), and this decision will be implemented for real payment.

Thus, you should take each decision independently of all other decisions as if it's the one that counts, because it may be.

Next

Comprehension questions

The questions below test your understanding of the instructions.

Important: If you fail to answer any one of these questions correctly, you will not be allowed to participate in the rest of the study, and you will not be able to earn a bonus.

1. Which one of the following statements is correct if the following lottery is played for you?

   With probability **60%: Get $ 15**
   With probability **40%: Get $ 5**

   Please select one of the statements:

   - ○ It is possible that I get paid both $15 and $5, i.e., I may receive a total amount of $20 from this lottery.
   - ○ I receive EITHER $15 OR $5 from this lottery.
   - ○ It is possible that I receive no money from this lottery.

2. Suppose a person made the decisions shown in the picture below. Which of the following statements is correct regarding these decisions?

Option A			Option B
	●	○	With certainty: Get $ 0
	●	○	With certainty: Get $ 1
	●	○	With certainty: Get $ 2
	●	○	With certainty: Get $ 3
	●	○	With certainty: Get $ 4
	●	○	With certainty: Get $ 5
With probability **60%**: Get $ 15	●	○	With certainty: Get $ 6
	●	○	With certainty: Get $ 7
With probability **40%**: Get $ 5	●	○	With certainty: Get $ 8
	○	●	With certainty: Get $ 9
	○	●	With certainty: Get $ 10
	○	●	With certainty: Get $ 11
	○	●	With certainty: Get $ 12
	○	●	With certainty: Get $ 13
	○	●	With certainty: Get $ 14
	○	●	With certainty: Get $ 15

   Please select one of the statements:

   - ○ This person indicated that the lottery is worth more to them than $9.
   - ○ This person indicated that the lottery is worth between $3 and $7 to them.
   - ○ This person indicated that the lottery is worth between $8 and $9 to them.

3. Now imagine that the person who made the decisions above is uncertain about how much exactly the lottery is worth to them.
   This person is certain, however, that the lottery is worth between $6 and $11 to them. Please position the slider to accurately reflect this level of certainty:

   very uncertain **completely certain**

   This person is certain that the lottery is worth **Use the slider!** to them.

Next

L.2 Treatment *Low Default Risk*

Welcome

Thank you for participating in this study.

This study will take approximately 25 minutes to complete.

You will earn a fixed reward of $0.50. In addition, you can earn a bonus for completing the study. To complete the study, you will need to read all instructions carefully and answer the corresponding comprehension questions correctly. You receive a **minimum bonus of $1.20** when you complete the study.

You can earn an ***additional* bonus**. At the end of the study, one of the tasks will be randomly selected and your decision in this task determines your additional bonus. The chance that you get paid an additional bonus for the first part is 1 in 3.

Important information

- You should think about each task **independently** of all other tasks in this study. There is no point in strategizing across tasks.
- You will note that we sometimes ask you to work on similar-sounding tasks. Theses tasks might have similar answers, or very different ones. Please consider each individual task **carefully**.

Next

Part 1: Instructions

Please read these instructions carefully. There will be comprehension checks. If you fail these, you will be excluded from the study and only receive the reward of $0.50.

In this study, there are various lotteries, all of which pay different amounts of money with different probabilities. To illustrate these probabilities, we will use the metaphor of colored balls in a bag. Imagine that there is a bag that contains 100 balls. Each of these balls has one of the following ten colors:

- red
- blue
- green
- orange
- brown
- pink
- black
- gold
- gray
- purple

The computer selects one of the 100 balls at random, where each ball is equally likely to get selected. Across lotteries, the number of balls of a given color might vary. Each color is associated with its own corresponding payout for you. An example lottery is:

70 out of 100 balls are **red**.
If a **red** ball is drawn: **Get $ 20**

3 out of 100 balls are **blue**.
If a **blue** ball is drawn: **Get $ 5**

3 out of 100 balls are **green**.
If a **green** ball is drawn: **Get $ 5**

3 out of 100 balls are **orange**.
If a **orange** ball is drawn: **Get $ 5**

3 out of 100 balls are **brown**.
If a **brown** ball is drawn: **Get $ 5**

3 out of 100 balls are **pink**.
If a **pink** ball is drawn: **Get $ 5**

3 out of 100 balls are **black**.
If a **black** ball is drawn: **Get $ 5**

3 out of 100 balls are **gold**.
If a **gold** ball is drawn: **Get $ 5**

3 out of 100 balls are **gray**.
If a **gray** ball is drawn: **Get $ 5**

6 out of 100 balls are **purple**.
If a **purple** ball is drawn: **Get $ 5**

You will always know how many balls of a given color are contained in the bag before making your decision. The lotteries will actually be played out by the computer and determine your earnings in exactly the way we describe on the following page.

In total, we have created 6 lotteries. For each lottery, you will make decisions on two consecutive decision screens.

- Decision screen 1: You will make various choices to indicate **how much** a lottery is worth to you.
- Decision screen 2: You will indicate **how certain** you are about how much **exactly** a lottery is worth to you.

Throughout the experiment, there are no right or wrong answers, because how much you like a lottery depends on your personal taste. We are only interested in learning about what you prefer.

Decision screen 1

On decision screen 1, you will be asked to choose which of two payment options you prefer. You will see choice lists such as the one below, where each row is a separate choice. In every list, the left-hand side option (Option A) is one of the lotteries that we generated, which is identical in all rows. Here, you can see how many balls of each color are contained in the bag that determines your payment. The right-hand side option (Option B) is a safe payment that you would receive *with certainty*, so there is no risk attached. This safe payment in Option B increases as you go down the list.

You should consider the choice in each row **independently** of all other rows, because in the end one row will be randomly selected for payout and you will receive the option that you selected in that row. In some choice lists that you complete, you will receive a budget. If so, this budget will be listed at the top of your screen and paid out to you along with the payouts from the lottery.

Usually, people start by preferring Option A for small certain amounts (at the top of the list). At some point they switch to Option B as they proceed down the list, because the certain amount associated with Option B increases, so that Option B becomes more attractive. **Thus, an effective way to complete these choice lists is to determine in which row you would like to switch from Option A to Option B.**

Based on your decisions in this choice list, we assess **how much the lottery is worth to you by using those certain payments where you switch from Option A to Option B.** For example, in the example choice list below, the lottery would be worth between $13 and $14 to you, because this is where switching occurs.

Option A			Option B
70 out of 100 balls are **red**. If a **red** ball is drawn: **Get $ 20** **3** out of 100 balls are **blue**. If a **blue** ball is drawn: **Get $ 5** **3** out of 100 balls are **green**. If a **green** ball is drawn: **Get $ 5** **3** out of 100 balls are **orange**. If a **orange** ball is drawn: **Get $ 5** **3** out of 100 balls are **brown**. If a **brown** ball is drawn: **Get $ 5** **3** out of 100 balls are pink. If a pink ball is drawn: **Get $ 5** **3** out of 100 balls are **black**. If a **black** ball is drawn: **Get $ 5** **3** out of 100 balls are gold. If a gold ball is drawn: **Get $ 5** **3** out of 100 balls are gray. If a gray ball is drawn: **Get $ 5** **6** out of 100 balls are **purple**. If a **purple** ball is drawn: **Get $ 5**	●	○	With certainty: **Get $ 0**
	●	○	With certainty: **Get $ 1**
	●	○	With certainty: **Get $ 2**
	●	○	With certainty: **Get $ 3**
	●	○	With certainty: **Get $ 4**
	●	○	With certainty: **Get $ 5**
	●	○	With certainty: **Get $ 6**
	●	○	With certainty: **Get $ 7**
	●	○	With certainty: **Get $ 8**
	●	○	With certainty: **Get $ 9**
	●	○	With certainty: **Get $ 10**
	●	○	With certainty: **Get $ 11**
	●	○	With certainty: **Get $ 12**
	●	○	With certainty: **Get $ 13**
	○	●	With certainty: **Get $ 14**
	○	●	With certainty: **Get $ 15**
	○	●	With certainty: **Get $ 16**
	○	●	With certainty: **Get $ 17**
	○	●	With certainty: **Get $ 18**
	○	●	With certainty: **Get $ 19**
	○	●	With certainty: **Get $ 20**

Click "Next" to read about decision screen 2.

Next

Decision screen 2 (for the same lottery)

For any given lottery, you may actually be **uncertain about how much money it is really worth** to you. For some lotteries, you may know exactly how much they are worth to you. For other lotteries, you may feel uncertain about whether a lottery is worth, say, $12, $13, $14, $15, or $16 to you.

On the second decision screen for a given lottery, we will hence ask you to use a slider to indicate **how certain** you are that your decisions on the first decision screen correspond **exactly** to how much the lottery is **worth to you**.

- The slider is linked to the statement below the slider. As you move around the slider, the values in this statement change. Your task is to position the slider in such a way that the statement corresponds to how you feel about the lottery.
- The **further you move the slider to the right, the more certain** you are about how much exactly the lottery is worth to you.
- For every step that you move the slider **further to the left (the less certain you are)**, the range of dollar values that you consider possible increases by $0.50.

Example

Suppose that on the first decision screen you indicated that the lottery below is worth between $13 and $14 to you (because you switched from Option A to B between $13 and $14). Your screen would then look like this:

Example lottery:

70 out of 100 balls are **red**.
If a **red** ball is drawn: **Get $ 20**

3 out of 100 balls are **blue**.
If a **blue** ball is drawn: **Get $ 5**

3 out of 100 balls are **green**.
If a **green** ball is drawn: **Get $ 5**

3 out of 100 balls are **orange**.
If a **orange** ball is drawn: **Get $ 5**

3 out of 100 balls are **brown**.
If a **brown** ball is drawn: **Get $ 5**

3 out of 100 balls are **pink**.
If a **pink** ball is drawn: **Get $ 5**

3 out of 100 balls are **black**.
If a **black** ball is drawn: **Get $ 5**

3 out of 100 balls are **gold**.
If a **gold** ball is drawn: **Get $ 5**

3 out of 100 balls are **gray**.
If a **gray** ball is drawn: **Get $ 5**

6 out of 100 balls are **purple**.
If a **purple** ball is drawn: **Get $ 5**

How certain are you that to you this lottery **is worth *exactly* the same as getting between $13 and $14** for sure?

very uncertain **completely certain**

I am certain that the lottery is worth **Use the slider!** to me.

Next

Your payment for part 1

If this part is randomly selected for payment in the end, your *additional* bonus will be determined as follows:

The computer will randomly select one choice list for payout, with equal probability. We then randomly select one of your decisions from this choice list (with equal probability), and this decision will be implemented for real payment.

Thus, you should take each decision independently of all other decisions as if it's the one that counts, because it may be.

Next

Comprehension questions

The questions below test your understanding of the instructions.

Important: If you fail to answer any one of these questions correctly, you will not be allowed to participate in the rest of the study, and you will not be able to earn a bonus.

1. Which one of the following statements is correct if the following lottery is played for you?

60 out of 100 balls are **red**.
If a **red** ball is drawn: **Get $ 15**

4 out of 100 balls are **blue**.
If a **blue** ball is drawn: **Get $ 5**

4 out of 100 balls are **green**.
If a **green** ball is drawn: **Get $ 5**

4 out of 100 balls are **orange**.
If a **orange** ball is drawn: **Get $ 5**

4 out of 100 balls are **brown**.
If a **brown** ball is drawn: **Get $ 5**

4 out of 100 balls are **pink**.
If a **pink** ball is drawn: **Get $ 5**

4 out of 100 balls are **black**.
If a **black** ball is drawn: **Get $ 5**

4 out of 100 balls are **gold**.
If a **gold** ball is drawn: **Get $ 5**

4 out of 100 balls are **gray**.
If a **gray** ball is drawn: **Get $ 5**

8 out of 100 balls are **purple**.
If a **purple** ball is drawn: **Get $ 5**

Please select one of the statements:

- ◯ It is possible that I get paid both $15 and $5, i.e., I may receive a total amount of $20 from this lottery.
- ◯ Only one of the colors will be drawn, hence I receive EITHER $15 OR $5 from this lottery.
- ◯ It is possible that I receive no money from this lottery.

2. Suppose a person made the decisions shown in the picture below. Which of the following statements is correct regarding these decisions?

Option A			Option B
60 out of 100 balls are **red**. If a **red** ball is drawn: **Get $ 15** **4** out of 100 balls are **blue**. If a **blue** ball is drawn: **Get $ 5** **4** out of 100 balls are **green**. If a **green** ball is drawn: **Get $ 5** **4** out of 100 balls are **orange**. If a **orange** ball is drawn: **Get $ 5** **4** out of 100 balls are **brown**. If a **brown** ball is drawn: **Get $ 5** **4** out of 100 balls are **pink**. If a **pink** ball is drawn: **Get $ 5** **4** out of 100 balls are **black**. If a **black** ball is drawn: **Get $ 5** **4** out of 100 balls are **gold**. If a **gold** ball is drawn: **Get $ 5** **4** out of 100 balls are gray. If a gray ball is drawn: **Get $ 5** **8** out of 100 balls are **purple**. If a **purple** ball is drawn: **Get $ 5**	◉	◯	With certainty: **Get $ 0**
	◉	◯	With certainty: **Get $ 1**
	◉	◯	With certainty: **Get $ 2**
	◉	◯	With certainty: **Get $ 3**
	◉	◯	With certainty: **Get $ 4**
	◉	◯	With certainty: **Get $ 5**
	◉	◯	With certainty: **Get $ 6**
	◉	◯	With certainty: **Get $ 7**
	◉	◯	With certainty: **Get $ 8**
	◯	◉	With certainty: **Get $ 9**
	◯	◉	With certainty: **Get $ 10**
	◯	◉	With certainty: **Get $ 11**
	◯	◉	With certainty: **Get $ 12**
	◯	◉	With certainty: **Get $ 13**
	◯	◉	With certainty: **Get $ 14**
	◯	◉	With certainty: **Get $ 15**

Please select one of the statements:

- ◯ This person indicated that the lottery is worth more to them than $9.
- ◯ This person indicated that the lottery is worth between $3 and $7 to them.
- ◯ This person indicated that the lottery is worth between $8 and $9 to them.

3. Now imagine that the person who made the decisions above is uncertain about how much exactly the lottery is worth to them.
   This person is certain, however, that the lottery is worth between $6 and $11 to them. Please position the slider to accurately reflect this level of certainty:

very uncertain **completely certain**

This person is certain that the lottery is worth **Use the slider!** to them.

Next

L.3 Treatment *Baseline Beliefs*

Welcome

Thank you for participating in this study.

This study will take approximately 25 minutes to complete.

You will earn a fixed reward of $0.50. In addition, you can earn a bonus if you complete the study. To complete the study, you will need to read all instructions carefully and answer the corresponding comprehension questions correctly. You receive a minimum bonus of $1.20 when you complete the study.

You can earn an ***additional*** **bonus**. At the end of the study, one of the tasks will be randomly selected and your decision in this task determines your additional bonus. The chance that you get paid an additional bonus for the first part is 1 in 3.

Important information

- You should think about each task **independently** of all other tasks in this study. There is no point in strategizing across tasks.
- You will note that we sometimes ask you to work on similar-sounding tasks. Theses tasks might have similar answers, or very different ones. Please consider each individual task **carefully**.
- Whenever a task involves a random draw, then this random draw will **actually be implemented for you by the computer** in exactly the way it is described to you in the task.

Next

Part 1: The guessing tasks

Please read these instructions carefully. We will test your understanding of them later.

In this study, you will be asked to complete **6 guessing tasks**.
In each guessing task, there are two bags, "bag A" and "bag B". Each bag contains 100 balls, some of which are red and some of which are blue. One of the bags is selected at random by the computer as described below. You will not observe which bag was selected. Instead, the computer will then randomly draw one or several balls from the secretly selected bag, and will show these balls to you. Your task is to **guess which bag was selected** based on the available information. The exact procedure is described below:

Task setup

- There is a deck of cards that consists of 100 cards. Each card in the deck either has "bag A" or "bag B" written on it. You will be informed about **how many** of these 100 cards have "bag A" or "bag B" written on them.
- There are two bags, "bag A" and "bag B". You will be informed about **how many red and blue balls** each bag contains.

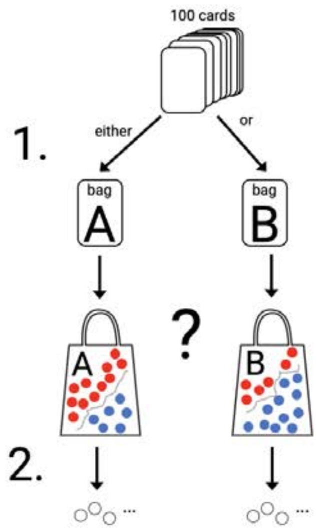

Sequence of events

1. The computer **randomly selects one** of the 100 cards, with equal probability.
   If a "bag A" card was drawn, bag A is selected.
   If a "bag B" card was drawn, bag B is selected.
2. Next, the computer **randomly draws one or more** balls from the secretly selected bag. Each ball is equally likely to get selected. Importantly, if more than one ball is drawn, the computer draws these balls **with replacement**. This means that after a ball has been drawn and taken out of the bag, it gets replaced by a ball of the same color. Thus, the probability of a red or blue ball being drawn does not depend on whether previous draws were red or blue.
   The computer will **inform you about the color** of the randomly drawn balls.

You then make your guess by **stating a probability between 0% and 100%** that bag A was drawn. The corresponding probability that bag B was drawn is 100 minus your stated probability that bag A was drawn.

Please note:

- The number of "bag A" and "bag B" cards **varies across the guessing tasks**.
- The number of red and blue balls in each bag **varies across the guessing tasks**.
- The computer **draws a new card in each task**, so you should **think about which bag was selected in a task independently** of all other tasks.

Next

Your payment for part 1

There is a **prize of $10.00**. Whether or not you receive the $10.00 depends on how much probability you assigned to the bag that was *actually drawn* in that problem.

This means: if bag A was selected, your chances of receiving $10.00 are greater the higher the probability you assigned to bag A. If bag A was not selected, your chances of receiving $10.00 are greater the lower the probability you assigned to bag A. In case you're interested, the specific method that determines whether you get the prize is explained below:

A number q between 0 and 2500 is randomly drawn by the computer.
If bag A was selected in that problem, you receive $10.00 if the square of the probability (in percent) that you assigned to bag B is *lower than q*.
If bag A was not selected in that problem, you receive $10.00 if the square of the probability (in percent) that you assigned to bag A is *lower than q*.

All this means that, in order to earn as much money as possible, you should try to give your best estimate of the probability that each bag was drawn. For example, if you are 80% sure that bag A was selected and 20% sure that bag B was selected, you should allocate probability 80% to bag A and 20% to bag B.

Next

Your **certainty** about your guess

The optimal guess

Using the laws of probability, the computer computes a **statistically correct statement of the probability that bag A was drawn**, based on all the information available to you. This **optimal guess** does not rely on information that you do not have. It is just the best possible (this means: payoff-maximizing) estimate given the available information. In technical terms, this guess is based on a statistical rule called Bayes' Law.

Your certainty about your guess

In any given task, you may actually be **uncertain about whether your probability guess corresponds to the optimal guess**. On a separate decision screen, we will hence ask you to use a slider to indicate **how certain** you are that your guess equals the optimal guess.

- The slider is linked to the statement below the slider. As you move around the slider, the values in this statement change. Your task is to position the slider in such a way that the statement corresponds to how close you think your guess is to the optimal guess.
- The further you move the slider to the right, the more certain you are that your guess is close to the optimal guess.
- For every step that you move the slider further to the left (the less certain you are), the range of probability values that you consider possible increases by 2.

Example

Suppose that you stated a guess of 80%. Your screen would then look like this:

How **certain** are you that **the *optimal* guess is *exactly* 80%**?

Use the slider to complete the statement below.

very uncertain **completely certain**

I am certain that the optimal guess of the **probability that bag A** was drawn is **Use the slider!**.

Replacing your guess by the optimal guess

In each task, if you are **uncertain about what you should guess**, you may **increase your chances of winning** the prize of $10.00 by paying money to **replace your guess with the optimal guess**.

For this purpose, in each task, you receive a budget of $3.00. You then have to state the **highest amount (between $0.00 and $3.00) that you are willing to pay to replace your guess with the optimal guess.** In the end, a price p between $0.00 and $3.00 will be randomly determined by the computer. You will purchase the optimal guess at price *p* if *p* is below your stated amount, and you will not purchase the optimal guess and keep your endowment otherwise. If you buy the optimal guess, your own guess is replaced with the optimal guess with 10% probability. You only have to pay for the optimal guess if your guess actually gets replaced.

Next

Sequence of events in each task

You will be asked to complete 6 guessing tasks. For each task, there will be **3 decision screens:**

Decision screen 1

You will be asked to state the highest amount that you are **willing to pay** to **replace your own guess** (that you will make on decision screen 2) with the **<u>optimal guess</u> in this task**.

Decision screen 2

You have to guess which bag was selected by entering a probability (between 0 and 100) that expresses **how likely** you think it is that bag A as opposed to bag B have been selected.

Decision screen 3

You will be asked to indicate **how certain** you are that the guess you provided on decision screen 1 **equals the optimal guess in this task**.

Next

Comprehension questions

The following questions test your understanding of the instructions.

Important: If you fail to answer any one of these questions correctly, you will not be allowed to participate in the rest of the study, and you will not be able to earn a bonus.

Which statement about the number of cards corresponding to each bag is correct?

- ◯ The number of "bag A" cards is the same in all tasks.
- ◯ The exact number of cards corresponding to each bag may vary across tasks.

Which statement about the allocation of red and blue balls in the bags is correct?

- ◯ The exact fractions of red and blue balls in each bag may vary across tasks.
- ◯ The fraction of red balls in each bag is the same in all tasks, and bag A always contains the most red balls.

Which statement about your bonus payment is correct?

- ◯ Purchasing the optimal guess always means that I would make a loss.
- ◯ By purchasing the optimal guess I can potentially increase my earnings in the guessing tasks if my own guess is not sufficiently good.

In order to maximize your overall profit, how should you determine how much you are willing to pay for the optimal guess?

- ◯ I should be willing to pay a high amount if I think that my own guess will probably not be optimal.
- ◯ I should be willing to pay a high amount if I think that my own guess will probably be optimal.
- ◯ I should never pay money for the optimal guess because it costs money but has no benefits.

Next

L.4 Treatment *Low Default Beliefs*

Welcome

Thank you for participating in this study.

This study will take approximately 25 minutes to complete.

You will earn a fixed reward of $0.50. In addition, you can earn a bonus if you complete the study. To complete the study, you will need to read all instructions carefully and answer the corresponding comprehension questions correctly. You receive a minimum bonus of $1.20 when you complete the study.

You can earn an ***additional* bonus**. At the end of the study, one of the tasks will be randomly selected and your decision in this task determines your additional bonus. The chance that you get paid an additional bonus for the first part is 1 in 3.

Important information

- You should think about each task **independently** of all other tasks in this study. There is no point in strategizing across tasks.
- You will note that we sometimes ask you to work on similar-sounding tasks. Theses tasks might have similar answers, or very different ones. Please consider each individual task **carefully**.
- Whenever a task involves a random draw, then this random draw will **actually be implemented for you by the computer** in exactly the way it is described to you in the task.

Next

Part 1: The guessing tasks

Please read these instructions carefully. We will test your understanding of them later.

In this study, you will be asked to complete **6 guessing tasks**.
In each guessing task, there are ten bags, "bag A", "bag B", "bag C", "bag D", "bag E", "bag F", "bag G", "bag H", "bag I" and "bag J". Each bag contains 100 balls, some of which are red and some of which are blue. One of the bags is selected at random by the computer as described below. You will not observe which bag was selected. Instead, the computer will then randomly draw one or several balls from the secretly selected bag, and will show these balls to you. Your task is to **guess which bag was selected** based on the available information. The exact procedure is described below:

Task setup

- There is a deck of cards that consists of 100 cards. Each card in the deck either has "bag A", "bag B", "bag C", "bag D", "bag E", "bag F", "bag G", "bag H", "bag I" or "bag J" written on it.
- You will be informed about **how many** of these 100 cards have "bag A", "bag B", ..., or "bag J" written on them.
- There are ten bags, "bag A" through "bag J". You will be informed about **how many red and blue balls** each bag contains.

Sequence of events

1. The computer **randomly selects one** of the 100 cards, with equal probability.
   If a "bag A" card was drawn, bag A is selected.
   If a "bag B" card was drawn, bag B is selected.
   ... etc. ...
2. Next, the computer **randomly draws one or more** balls from the secretly selected bag. Each ball is equally likely to get selected. Importantly, if more than one ball is drawn, the computer draws these balls **with replacement**. This means that after a ball has been drawn and taken out of the bag, it gets replaced by a ball of the same color. Thus, the probability of a red or blue ball being drawn does not depend on whether previous draws were red or blue.
   The computer will **inform you about the color** of the randomly drawn balls.

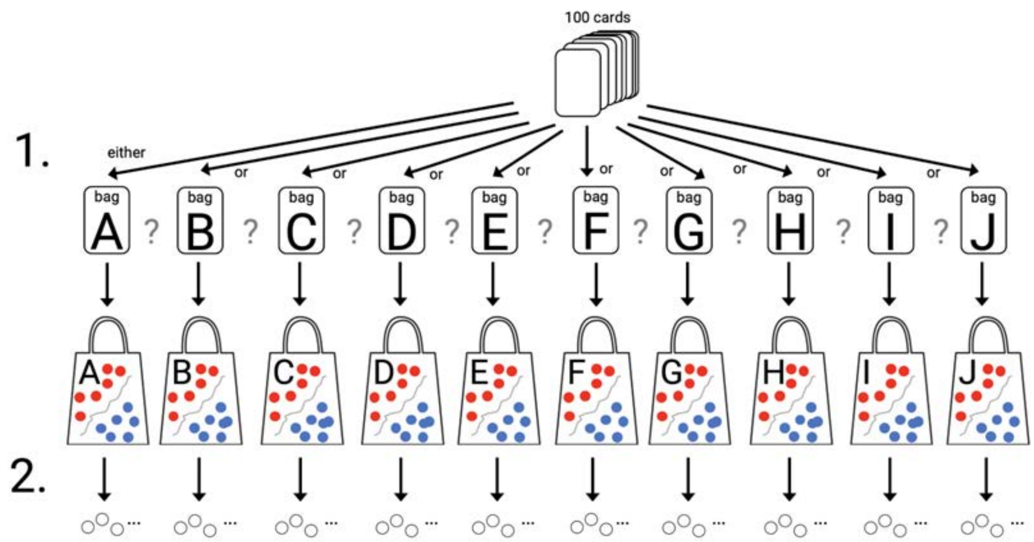

You then make your guess by **stating a probability between 0% and 100%** that bag A was drawn. The corresponding probability that one of bag B, C, D, E, F, G, H, I or J was drawn is 100 minus your stated probability that bag A was drawn.

Please note:

- The number of "bag A", "bag B", ..., and "bag J" cards **varies across the guessing tasks**.
- The number of red and blue balls in each bag **varies across the guessing tasks**.
- The computer **draws a new card in each task**, so you should **think about which bag was selected in a task independently** of all other tasks.

Next

Your payment for part 1

There is a **prize of $10.00**. Whether or not you receive the $10.00 depends on how much probability you assigned to the bag that was *actually drawn* in that problem.

This means: if bag A was selected, your chances of receiving $10.00 are greater the higher the probability you assigned to bag A. If bag A was not selected, your chances of receiving $10.00 are greater the lower the probability you assigned to bag A. In case you're interested, the specific method that determines whether you get the prize is explained below:

A number q between 0 and 2500 is randomly drawn by the computer.
If bag A was selected in that problem, you receive $10.00 if the square of the probability (in percent) that you assigned to all other bags is *lower than q*.
If bag A was not selected in that problem, you receive $10.00 if the square of the probability (in percent) that you assigned to bag A is *lower than q*.

All this means that, in order to earn as much money as possible, you should try to give your best estimate of the probability that each bag was drawn. For example, if you are 80% sure that bag A was selected and 20% sure that any bag other than bag A was selected, you should allocate probability 80% to bag A and 20% to bag B.

Next

Your **certainty** about your guess

The optimal guess

Using the laws of probability, the computer computes a **statistically correct statement of the probability that bag A was drawn**, based on all the information available to you. This **optimal guess** does not rely on information that you do not have. It is just the best possible (this means: payoff-maximizing) estimate given the available information. In technical terms, this guess is based on a statistical rule called Bayes' Law.

Your certainty about your guess

In any given task, you may actually be **uncertain about whether your probability guess corresponds to the optimal guess**. On a separate decision screen, we will hence ask you to use a slider to indicate **how certain** you are that your guess equals the optimal guess.

- The slider is linked to the statement below the slider. As you move around the slider, the values in this statement change. Your task is to position the slider in such a way that the statement corresponds to how close you think your guess is to the optimal guess.
- The further you move the slider to the right, the more certain you are that your guess is close to the optimal guess.
- For every step that you move the slider further to the left (the less certain you are), the range of probability values that you consider possible increases by 2.

Example

Suppose that you stated a guess of 80%. Your screen would then look like this:

How **certain** are you that **the *optimal* guess is *exactly* 80%**?

Use the slider to complete the statement below.

very uncertain **completely certain**

I am certain that the optimal guess of the **probability that bag A** was drawn is **Use the slider!**.

Replacing your guess by the optimal guess

In each task, if you are **uncertain about what you should guess**, you may **increase your chances of winning** the prize of $10.00 by paying money to **replace your guess with the optimal guess**.

For this purpose, in each task, you receive a budget of $3.00. You then have to state the **highest amount (between $0.00 and $3.00) that you are willing to pay to replace your guess with the optimal guess.** In the end, a price p between $0.00 and $3.00 will be randomly determined by the computer. You will purchase the optimal guess at price *p* if *p* is below your stated amount, and you will not purchase the optimal guess and keep your endowment otherwise. If you buy the optimal guess, your own guess is replaced with the optimal guess with 10% probability. You only have to pay for the optimal guess if your guess actually gets replaced.

Next

Sequence of events in each task

You will be asked to complete 6 guessing tasks. For each task, there will be **3 decision screens:**

Decision screen 1

You will be asked to state the highest amount that you are **willing to pay** to **replace your own guess** (that you will make on decision screen 2) with the **optimal guess in this task**.

Decision screen 2

You have to guess which bag was selected by entering a probability (between 0 and 100) that expresses **how likely** you think it is that bag A as opposed to any bag other than bag A have been selected.

Decision screen 3

You will be asked to indicate **how certain** you are that the guess you provided on decision screen 1 **equals the optimal guess in this task**.

Next

Comprehension questions

The following questions test your understanding of the instructions.

Important: If you fail to answer any one of these questions correctly, you will not be allowed to participate in the rest of the study, and you will not be able to earn a bonus.

Which statement about the number of cards corresponding to each bag is correct?

- ◯ The number of "bag A" cards is the same in all tasks.
- ◯ The exact number of cards corresponding to each bag may vary across tasks.

Which statement about the allocation of red and blue balls in the bags is correct?

- ◯ The exact fractions of red and blue balls in each bag may vary across tasks.
- ◯ The fraction of red balls in each bag is the same in all tasks, and bag A always contains the most red balls.

Which statement about your bonus payment is correct?

- ◯ Purchasing the optimal guess always means that I would make a loss.
- ◯ By purchasing the optimal guess I can potentially increase my earnings in the guessing tasks if my own guess is not sufficiently good.

In order to maximize your overall profit, how should you determine how much you are willing to pay for the optimal guess?

- ◯ I should be willing to pay a high amount if I think that my own guess will probably not be optimal.
- ◯ I should be willing to pay a high amount if I think that my own guess will probably be optimal.
- ◯ I should never pay money for the optimal guess because it costs money but has no benefits.

Next

L.5 Survey Expectations

Part 2 of this study

You have completed part 1. We will now continue with part 2 of this study.

Next

Your payment for part 2

In this part, there will be 3 tasks. At the end, one of the tasks will be randomly selected to count for your potential bonus. The chance that you get paid an additional bonus for this part is 1 in 3.

In each task, you will be asked to **state a guess in the form a probability estimate (between 0 and 100).**

There is a **prize of $2.00**. In each guessing task, there are two possible events, call them A and B. One of the two events actually occurred, the other did not. Whether or not you receive the $2.00 depends on **how much probability** you assigned to the **event that actually occurred**.

If event A occurred, your chances of receiving $2.00 are greater the higher the probability you assigned to event A. If event B occurred, your chances of receiving $2.00 are greater the higher the probability you assigned to event B.

In case you're interested, the specific method that determines whether you get the prize is explained below:

A number **q** between 0 and 2,500 is randomly drawn by the computer.

- If event A occurred in that problem, you receive $2.00 if the square of the probability (in percent) that you assigned to event B is *lower than q*.
- If event B occurred in that problem, you receive $2.00 if the square of the probability (in percent) that you assigned to event A is *lower than q*.

All this means that, in order to earn as much money as possible, you should try to provide your best probability estimate in each task.

Next

The study begins on the next page

We will now start with part 2 of study.

Next