Insensitive Investors

CONSTANTIN CHARLES, CARY FRYDMAN, and METE KILIC *

ABSTRACT

We experimentally study the transmission of subjective expectations into actions. Subjects in our experiment report valuations that are far too insensitive to their expectations, relative to the prediction from a frictionless model. We propose that the insensitivity is driven by a noisy cognitive process that prevents subjects from precisely computing asset valuations. The empirical link between subjective expectations and actions becomes stronger as subjective expectations approach rational expectations. Our results highlight the importance of incorporating weak transmission into belief-based asset pricing models. Finally, we discuss how cognitive noise can provide a microfoundation for inelastic demand in the stock market.

*Charles is at the London School of Economics. Frydman and Kilic are at the Marshall School of Business, University of Southern California. We are grateful to Stefan Nagel (editor), an associate editor, two anonymous referees, Marianne Andries, Elena Asparouhova, Nicholas Barberis, Sabine Bernard, Alex Chinco, Ricardo De La O, John Duffy, Xavier Gabaix, Stefano Giglio, Thomas Graeber, Sam Hartzmark, David Hirshleifer, Kristy Jansen, Lawrence Jin, Ralph Koijen, Stijn van Nieuwerburgh, Jeremy Stein, Selale Tuzel, Jessica Wachter, Sunil Wahal, Mitch Warachka as well as participants at the Miami Behavioral Finance Conference, NBER Asset Pricing Meeting, Colorado Finance Summit, California Corporate Finance Conference, Chapman University Experimental and Behavioral Finance Conference, Stanford Institute for Theoretical Economics, Chicago Booth Asset Pricing Conference, Middlebury Experimental Seminar, TBEAR Workshop, UC Irvine Finance Conference, Arizona State University, Brigham Young University, Purdue University, Southern Methodist University, Hong Kong University, City University of Hong Kong, and USC Marshall for helpful comments. Charles: c.charles@lse.ac.uk. Frydman: cfrydman@marshall.usc.edu. Kilic: mkilic@marshall.usc.edu. Frydman acknowledges NSF grant #1749824 for financial support. We have read the *Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Economists have spent the past several years using surveys to document facts about investors' expectations of stock returns. A clear fact that emerges from this literature is that the subjective expected returns that investors report on surveys systematically depart from objective expected returns (Greenwood and Shleifer (2014), Adam and Nagel (2023), Nagel and Xu (2023)). This fact rejects standard rational expectations models and has motivated a new class of asset pricing theories aimed at matching both subjective expectations and realized returns (e.g., Barberis et al. (2015), Hirshleifer, Li, and Yu (2015), Barberis et al. (2018), Bordalo et al. (2019), Jin and Sui (2022), Nagel and Xu (2022)). These models formalize the subjective expectation formation process in a psychologically grounded manner, but retain the standard assumption that investors fully act on their subjective expectations.

In a parallel strand of research, several authors have highlighted a puzzling disconnect between measured subjective expectations and investor actions. Using data from a sample of wealthy retail investors, Giglio et al. (2021a) document that the sensitivity of equity portfolio shares to subjective return expectations is an order of magnitude weaker than predicted by standard frictionless models. This weak transmission of beliefs to actions appears to be a robust phenomenon that is observed in a variety of other settings (Amromin and Sharpe (2014), Drerup, Enke, and Von Gaudecker (2017), Ameriks et al. (2020), Liu and Palmer (2021)). Even in times of a market crash, when investors arguably pay a lot of attention to the stock market, actions remain too insensitive to subjective beliefs (Giglio et al. (2021b)).[1]

In this paper, we conduct a series of controlled lab experiments to investigate why investor actions do not fully reflect the subjective beliefs reported on surveys. We test the hypothesis that the insensitivity between beliefs and actions is driven by a noisy cognitive process that prevents investors from precisely forming asset valuations, conditional on reported beliefs.

[1]The subjective beliefs reported on surveys clearly contain valuable information about portfolio choice and aggregate outcomes, such as fund flows (Greenwood and Shleifer (2014)). The puzzle raised in the recent literature is about the *quantitative* strength of the relationship between beliefs and actions.

Our hypothesis is motivated by a recent agenda in behavioral economics which argues that the decision-making process is subject to inherent *cognitive noise* (see Woodford (2020) for a recent review). The noise arises inside the investor's mind and leads to systematic decision biases. In particular, the investor is aware of this noise, and as a consequence, she optimally shades her decision towards a "cognitive default" value.

To build intuition for our proposed cognitive noise mechanism, consider an investor who is estimating the value of a risky asset. The investor will naturally lean on her beliefs about the asset's future payoffs, but she may find it difficult to compute her exact valuation given these beliefs and her risk appetite. We model this difficulty by assuming the investor only has access to a noisy cognitive signal about her true valuation. The noise could reflect uncertainty about beliefs, uncertainty about her own risk aversion, or the noisy cognitive process of translating beliefs and preferences into an asset valuation.

We further assume the investor is aware of the noise in her decision process, and she therefore optimally compresses her valuation towards a cognitive default. The cognitive default is the value that the investor would report *before* drawing her noisy cognitive signal. The cognitive default could, for example, be driven by an investor's past experience with valuing similar assets. The compression of valuation towards a cognitive default may be interpreted as a rule of thumb, but it can also be microfounded by Bayesian updating in the presence of cognitive noise (Gabaix (2019)). For our purposes, the important implication of cognitive noise is that compressing valuation towards a cognitive default immediately dampens the transmission of stated beliefs to valuations.

In our experiments, we shut down all institutional frictions that operate in the field, and we still uncover a weak transmission from beliefs to actions. This empirical result suggests that a psychological mechanism is responsible for the weak transmission; we then provide two targeted tests of the cognitive noise hypothesis. First, we exogenously reduce the complexity of the decision problem, which has been shown to reduce cognitive noise in the valuation process (Enke and Graeber (2023)). We find that the reduction in cognitive noise causally and substantially increases the passthrough from beliefs to actions. Second, we find that subjects incorporate their cognitive default into reported valuations in a manner that is quantitatively consistent with theory.

Our experiments contain a variety of novel design features that enable us identify the weak transmission effect and probe its underlying mechanism. In our first experiment (Experiment 1), subjects are given the opportunity to invest in a risky asset that has a time-varying expected payoff. We elicit both subjective beliefs and valuations about the asset in an incentive compatible fashion. A frictionless model would predict that when a subject increases her reported expected payoff by one unit, this should translate into a one unit increase in her willingness to pay (WTP). Our experimental data strongly depart from this frictionless benchmark: we find that a one unit increase in subjective expected payoff leads to only a 61% increase in WTP. Thus, we recover the weak transmission of beliefs to actions, even in the absence of any institutional frictions. Moreover, because we ask subjects for their WTP and beliefs on the same experimental screen, beliefs should be readily accessible, which arguably tilts the scales away from finding the weak transmission effect.

While the average subject in our first experiment does exhibit a weak transmission of beliefs to actions, we find substantial heterogeneity on this dimension. This motivates an important question: which subjects transmit their beliefs more vigorously into actions? We find that subjects whose beliefs are closer to the rational (Bayesian) benchmark are the ones who transmit their beliefs more strongly into valuations. Put differently, beliefs that are farther from the rational benchmark are less likely to be incorporated into valuations – though this is only correlational evidence.

In our second experiment (Experiment 2), we test whether cognitive noise causally affects the passthrough from beliefs to actions. To do so, we draw on the finding from Enke and Graeber (2023) that subjects report higher levels of cognitive noise when decisions are more complex. We argue that it is more complex to price an asset based on subjective beliefs that are learned from past realized payoffs compared to objective beliefs that are endowed. We therefore manipulate cognitive noise by varying whether beliefs are subjective or objective, but we hold constant the beliefs themselves.

We implement the cognitive noise manipulation through a novel design feature. We endow subjects in Experiment 2 with the beliefs reported by subjects from Experiment 1. Specifically, each subject in Experiment 2 is endowed with an objective payoff distribution, and we generate this payoff distribution from the subjective beliefs of a randomly matched partner in Experiment 1. To help convey the critical design aspect, suppose that after observing a sequence of realized payoffs, a subject from Experiment 1 reports a distribution of beliefs denoted by b_1 (and her associated WTP given these beliefs). In Experiment 2, there is no learning and we instead endow the subject with beliefs b_1 and ask her to price the asset conditional on these *objective* beliefs. Our manipulation is grounded in the hypothesis that cognitive noise is larger in settings where additional cognitive operations are needed, such as learning from past data and forming subjective beliefs (Findling and Wyart (2021)). By comparing the sensitivity of WTP to beliefs across Experiments 1 and 2, we can estimate the causal effect of cognitive noise.

We find that endowing subjects with objective beliefs leads to a striking difference in pricing behavior: for every unit increase in expected payoff, subjects in Experiment 2 increase their WTP by 82%, compared to only 61% in Experiment 1. At the same time, the transmission strength in Experiment 2 remains significantly below the frictionless benchmark, suggesting that cognitive noise from other sources besides belief formation affect valuation. Because we hold beliefs constant across experiments, our interpretation is that cognitive noise causally decreases the sensitivity of actions to beliefs. Moreover, to our knowledge, this is the first piece of evidence indicating that valuation is substantially less sensitive to subjective beliefs compared to objective beliefs.[2]

In our final experiment (Experiment 3), we test the core feature of the cognitive noise mechanism. The key to generating weak transmission under this mechanism is that valuation is compressed towards a cognitive default parameter. We develop a simple method for eliciting the cognitive default in each period: we ask subjects to report their valuation of the asset *before* observing the asset's objective payoff distribution. We interpret this valuation as the subject's cognitive default, since it represents their valuation before observing the objective payoff distribution, and thus before drawing their noisy cognitive signal. We then additionally ask for the subject's valuation after we present the objective payoff distribution; this is the valuation that represents the subject's WTP for the asset. Not only do we find that the cognitive default is correlated with WTP, but subjects apply a decision weight to their cognitive default that is quantitatively consistent with theory.

Overall, our experimental findings provide important guidance for the role of subjective expectations data in asset pricing. Brunnermeier et al. (2021) point to the need for more research on the interaction between beliefs and actions to better understand the role of expectations data for asset pricing. Our work highlights that the weak transmission of reported beliefs to valuations can arise in a simple environment that is insulated from institutional frictions. One lesson from our data is that researchers should be cautious when analyzing the quantitative predictions of models that assume investors fully act on their reported beliefs.

[2]Our finding is similar to, but distinct from, the experimental result in Hartzmark, Hirshman, and Imas (2021) where subjects react more strongly to information about goods that they own compared to those that they do not own. In Hartzmark, Hirshman, and Imas (2021), the endowment of an asset is randomly varied across treatments. In our setting, it is the endowment of *beliefs* that varies across treatments, and we find that WTP reacts more strongly when beliefs are endowed rather than learned.

At the same time, incorporating a weak transmission channel into existing belief-based asset pricing models may provide an opportunity to further improve quantitative fits to the data.

In a related paper on weak transmission, Andries et al. (2022) conduct an experiment in which they vary the signal informativeness about future returns. When subjects perceive the signal to be less informative, subjective expectations deviate from the Bayesian benchmark and allocations underreact to beliefs. To the extent that subjects in our experiment perceive objective beliefs to be more informative than subjective beliefs, our results are consistent with those of Andries et al. (2022). In contrast to their paper, a major focus of our analysis is to uncover the psychological mechanism that generates weak transmission. Our proposed cognitive noise mechanism may explain the underreaction of actions to beliefs observed by Andries et al. (2022) in a different experimental paradigm.

The rest of this paper is organized as follows. Section I presents a conceptual framework that illustrates how cognitive noise can generate the weak transmission of beliefs to WTP. Sections II - IV present the main results from our three experiments. In Section V we discuss alternative mechanisms for weak transmission and the broader implications of our results for asset pricing. Section VI concludes with directions for future work.

I. Conceptual framework

We begin by stating the relationship between beliefs and WTP under the frictionless benchmark. We then introduce our key assumption of cognitive noise, and derive its implications for valuation.

A. Frictionless benchmark

Suppose that at time t an agent can invest in an asset which delivers a stochastic payoff D_{t+1}. The agent forms beliefs about the payoff's conditional distribution, where the mean of this subjective distribution is given by $\mathbb{E}_t^*[D_{t+1}]$. After forming expectations, and before the payoff D_{t+1} is realized, the agent decides what price P_t she is willing to pay for a claim on D_{t+1}.[3] The agent's subjective expected return is therefore given by $\mathbb{E}_t^*[R_{t+1}] = \mathbb{E}_t^*[D_{t+1}]/P_t$. We can rewrite this identity as

$$r_t = d_t - p_t, \tag{1}$$

where $r_t = \log \mathbb{E}_t^*[R_{t+1}]$, $d_t = \log \mathbb{E}_t^*[D_{t+1}]$, and $p_t = \log P_t$. Unless otherwise noted, throughout the rest of this section we use WTP, subjective expected returns, and subjective expected payoffs in logs which will simplify the predictions that we derive here and test in the next three sections.

The subjective expected return r_t is equivalent to the discount rate that the agent applies to the subjective expected payoff d_t, in order to generate her WTP, p_t. This equivalence is easily seen by rearranging (1) into:

$$p_t = d_t - r_t. \tag{2}$$

Equation (2) implies that a one unit increase in the subjective expected payoff should translate into a one unit increase in WTP – controlling for her discount rate. This 1-1 relationship between beliefs and valuations will serve as the frictionless benchmark across all three of our experiments. We assume that the agent applies the following discount to the subjective expected payoff:

$$r_t = \gamma \lambda_t, \tag{3}$$

where γ is the price of risk (e.g., risk aversion) and λ_t is the quantity of risk implied by the

[3] One can think of such an asset as a dividend strip. The one-period nature of the asset simplifies the expectation formation process and is sufficient to convey our main conceptual insight.

agent's subjective beliefs (e.g., conditional volatility).[4]

B. Insensitive actions

We now consider a friction in the transmission of the agent's reported beliefs to her actions. The friction is motivated by a recent agenda in behavioral economics which argues that cognitive noise corrupts the decision-making process and leads to systematic biases (Woodford (2020)). As we will see, the main implication of this friction is that it induces an elasticity between beliefs and actions that is less than the frictionless benchmark of 1.

We define the agent's true valuation of the asset at time t as $p_t^* = d_t - \gamma\lambda_t$, where d_t and λ_t are the agent's reported expected payoff and perceived risk, respectively. The variable p_t^* is the benchmark price that is predicted by a frictionless model. Our key assumption is that cognitive noise prevents the agent from accessing this frictionless valuation due to cognitive and attentional constraints (Gabaix (2019), Enke and Graeber (2023)). Instead, she only has access to a noisy cognitive signal $p_t^0 = p_t^* + \epsilon_t = d_t - \gamma\lambda_t + \epsilon_t$, where ϵ_t is drawn from $N(0, \sigma_\epsilon^2)$. The investor herself generates the noisy cognitive signal when she is deliberating about her valuation.

In our setting, cognitive noise may be interpreted as difficulty with the process of valuing the asset conditional on beliefs, but it can also reflect uncertainty about valuation inputs such as beliefs or risk aversion. The agent exhibits less cognitive noise as she becomes more certain about her expectations. Yet, even when she is completely certain about her expectations, noise can still arise in the decision process that transforms precise beliefs into

[4]We assume that the time increment is short enough such that the riskless rate is zero and the discount rate only represents an instantaneous risk premium. We interpret this assumption as the agent perceiving the stochastic payoff as an instantaneous gamble with no need for time discounting, which will be the case in our experimental design.

actions.[5]

Following Gabaix (2019) and Enke and Graeber (2023) we adopt a Bayesian perspective whereby the agent has a prior over what her true valuation is: p_t^* is drawn from a normal distribution $N(\bar{p}_t, \sigma_p^2)$. Here $\bar{p}_t$ is a "cognitive default" parameter, which is the valuation she would report before drawing her noisy cognitive signal. We emphasize that $\bar{p}_t$ is a forward looking variable, in the sense that it represents the expected valuation that the agent holds before any further information about the asset is revealed. The agent then combines her prior and signal to arrive at the posterior mean, which is the WTP that she reports:

$$\begin{aligned} p_t &= (1-x)\bar{p}_t + x p_t^0 \\ &= (1-x)\bar{p}_t + x d_t - x\gamma\lambda_t + x\epsilon_t \end{aligned} \tag{4}$$

where $x = \sigma_p^2/(\sigma_p^2 + \sigma_\epsilon^2)$ is the weight she attaches to her noisy signal relative to the cognitive default.[6] It is important to highlight that $\bar{p}_t$ does not represent some "irrelevant anchor" that lowers the quality of the investor's decision-making. Rather, because cognitive noise prevents the investor from accessing her true valuation, she optimally incorporates information from her prior to arrive at her perceived valuation. See Appendix A for a derivation. The crucial implication encoded in Equation (4) is that a one unit increase in d_t now leads to an increase in p_t by only x units, where $0 < x < 1$.

[5] Our motivation for incorporating cognitive noise into a framework of asset valuation is partially driven by previous experimental work which documents sizeable and systematic effects of cognitive noise in even simpler environments (Woodford (2020)). For example, multiple studies have shown that experimental subjects make systematic errors when judging which of two symbolic numbers is larger (Moyer and Landauer (1967), Dehaene (2011), Frydman and Jin (2022)). Thus, given that cognitive noise operates in simple judgment tasks, we believe cognitive noise is likely to also be active in the more complex environment of asset valuation.

[6] While we assume that $\bar{p}_t$ can vary over time, we assume that σ_p^2 is constant. Together with the assumption that σ_ϵ^2 is constant, this implies that x does not vary over time. In Internet Appendix IA.1, we report a test of this implication, and our data support the idea that x does not vary over the course of each experiment.

To help illuminate the mapping between our framework and applications, consider an investor who is assessing her valuation of the aggregate stock market. If good fundamental news is released about the market, then the investor updates her beliefs about future cash flows, which corresponds to an increase in d_t in our framework. But it may be very difficult for the investor to figure out exactly how much this shift in cash flow expectations should shift her WTP for the stock market. We model this friction as the additive noise term, ϵ_t. The investor is aware of this difficulty, and when coming up with her WTP, she therefore leans on a default price; we interpret the default price as the expected price the agent would be willing to pay, without first incorporating the fundamental news. The default price is represented by $\bar{p}_t$ in our framework. If enough investors behave in this manner – and hold a similar default – then the price of the stock market will adjust in the right direction, but not by enough.

We note that Equation (4) also implies that WTP will sluggishly respond to perceived risk. In other words, the weak transmission of beliefs to valuation is not confined to the first moment of the subjective payoff distribution, but also operates over our assumed measure of perceived risk. This is a testable prediction that we will take to our experimental data later in the paper.

C. Predictions

Here we summarize the main implications of our conceptual framework and develop three testable predictions that will guide our experimental design. Our first prediction summarizes the basic implication of cognitive noise for valuation.

Prediction 1. *If the investor has cognitive noise (i.e., $\sigma_\epsilon^2 > 0$), then a 1 unit increase in the expected dividend will lead to an increase in willingness to pay for the asset of only $x < 1$ units.*

Prediction 1 indicates that the weak transmission from beliefs to actions can persist even in a controlled laboratory setting, where all plausible institutional reasons for the weak transmission are shut down by design. Our next prediction provides comparative statics about the level of cognitive noise and the insensitivity of actions to beliefs.

Prediction 2. *If cognitive noise is exogenously decreased (i.e., $\sigma_\epsilon^2 > 0$ is decreased) then the the passthrough from beliefs to actions should become stronger. That is, x should increase as we decrease cognitive noise.*

Finally, as a deeper test of the psychological mechanism, we rely on a quantitative prediction that is implied by Equation (4). Specifically, when forming asset valuations, the investor attaches a positive decision weight to both her cognitive default and her reported expectations, and these decision weights sum to 1.

Prediction 3. *Suppose the investor has cognitive noise (i.e., $\sigma_\epsilon^2 > 0$). When estimating Equation (4), the coefficients on d_t and $\bar{p}_t$ should (i) both be positive and (ii) sum to 1.*

Implementing a test of Prediction 3 requires observing $\bar{p}_t$; later in the paper we develop an experimental method for measuring the cognitive default. In the next three sections, we test our predictions in an experimental setting. Predictions 1, 2, and 3 will be tested in Experiments 1, 2, and 3, respectively.

II. Experiment 1: Identifying weak transmission in a controlled lab setting

A. Experimental design

The goal of our first experiment is to cleanly test for the weak transmission of beliefs to WTP. Importantly, our design shuts down several factors that can generate a low sensitivity of actions to beliefs in the field, such as capital gains taxes, default options in retirement plans, and costly portfolio monitoring. Additionally, and in contrast to standard survey methodologies, we incentivize the elicitation of beliefs and valuations.

In our design there is a stock that pays a dividend, D_t, in each of 30 periods. There are five possible values for the dividend: $\{\$60, \$85, \$115, \$135, \$150\}$. This five point distribution of payoffs is similar to the distribution of returns that Giglio et al. (2021a) use to elicit beliefs from their survey respondents.[7] The conditional distribution of D_t is governed by a two-state Markov chain. We denote the state in period t by s_t, which can take on one of two values, either *good* or *bad*. In the *bad* state, the distribution of dividends is given by:

$$\Pr(D_t|s_t = \textit{bad}) \equiv (\$60, 0.15; \$85, 0.30; \$115, 0.40; \$135, 0.10; \$150, 0.05) \tag{5}$$

In the *good* state, the distribution of dividends is given by:

$$\Pr(D_t|s_t = \textit{good}) \equiv (\$60, 0.05; \$85, 0.10; \$115, 0.40; \$135, 0.30; \$150, 0.15). \tag{6}$$

The distribution in the *good* state has a higher mean and lower volatility, compared with the distribution in the *bad* state. We initialize the state in period 1 to be either *good* or *bad* with equal probability: $\Pr(s_1 = \textit{good}) = 50\%$. The states are persistent; the probability of remaining in the same state from one period to the next is 80%. Therefore, with 20% probability, the state switches in each period.

Subjects are given all the above information about the model of dividends; however, they do not observe the identity of the state in each period. As such, subjects face a learning problem in which they can use data on past dividends to infer the probability that the current state is good. We choose the above stochastic process because it induces substantial time series variation in the expected dividend. Moreover, the two-state switching process

[7] Giglio et al. (2021a) elicit a distribution over five different ranges of returns, whereas we elicit a distribution over five different values of the dividend.

guarantees that the variation does not decline over time (as would be the case in, say, a model where the probability of switching from one state to the other is zero). For our purposes, the substantial time series variation in expected dividends is useful for estimating the passthrough from beliefs to WTP. While it is difficult to guarantee that *subjective* beliefs will inherit this same amount of time series variation, we reason that the two-state switching process gives us a good chance of observing large fluctuations in subjective beliefs. To ease comparability of behavior across subjects, we use the same realized sequence of thirty dividends for all subjects.[8]

In 8 randomly chosen periods, we elicit a subject's full distribution of beliefs about the next period's dividend. In the other 22 periods, we do not elicit beliefs, and subjects simply observe the realized dividend.[9] Specifically, we ask subjects for the probability that they attach to each of the five possible dividend outcomes. The ordering of the buckets (i.e., lowest to highest or highest to lowest) is randomized at the subject level, and we enforce that the probabilities add up to 100%. We also ask subjects to report the price they are willing to pay for the right to receive next period's dividend, D_{t+1}. These two elicitations enable us to test the relation between subjective payoff expectations and WTP as well as how the subjective payoff distribution differs from the objective payoff distribution.

Importantly, we incentivize the expectations question and the WTP question. When we elicit a subject's beliefs about next period's dividend, we pay subjects based on their accuracy relative to how a Bayesian agent would respond. To see how a Bayesian agent would respond, we derive the probability that the state is bad, conditional on all past dividends. We denote

[8] There is a tradeoff in using the same realized sequence of dividends for all subjects. On the one hand, the particular sequence may not be representative of the true data generating process. On the other hand, fixing the sequence for all subjects enables a more precise estimate of the cross-sectional variation in subjective beliefs and in WTP – which we explore more deeply in the next subsection.

[9] We elicit beliefs in the same 8 (randomly chosen) periods for all subjects. See Internet Appendix IA.2 for screenshots of the experiment.

this probability as $q_t = \Pr(s_t = bad|D_t, D_{t-1}, ..., D_1)$. Conditional on q_t, the distribution of dividends can be computed using the distributions in the *good* and *bad* states shown in Equations (5) and (6). Because the stochastic process is Markovian, we can rewrite the expression for q_t as a function of the current period's realized dividend and the prior belief:

$$\begin{aligned} & q_t(q_{t-1}, D_t) \\ & = \frac{\Pr(D_t|s_t = bad)\Pr(s_t = bad|q_{t-1})}{\Pr(D_t|s_t = bad)\Pr(s_t = bad|q_{t-1}) + \Pr(D_t|s_t = good)\Pr(s_t = good|q_{t-1})} \\ & = \frac{\Pr(D_t|s_t = bad)(0.8q_{t-1} + 0.2(1 - q_{t-1}))}{\Pr(D_t|s_t = bad)(0.8q_{t-1} + 0.2(1 - q_{t-1})) + \Pr(D_t|s_t = good)(0.2q_{t-1} + 0.8(1 - q_{t-1}))}, \end{aligned} \tag{7}$$

where the expressions $\Pr(D_t|s_t = bad)$ and $\Pr(D_t|s_t = good)$ are defined in Equations (5) and (6) (Frydman et al. (2014)). Given the probability that the stock is in the bad state, the expected dividend is just a weighted average of the expected dividend in each of the two states: $\mathbb{E}[D_{t+1}|q_t] = q_t\mathbb{E}[D_{t+1}|s_t = bad] + (1-q_t)\mathbb{E}[D_{t+1}|s_t = good]$. Similarly, the probability of each dividend outcome is a weighted average of the probability of that outcome in each of the two states. For example, for a \$60 dividend, $\Pr(D_{t+1} = \$60|q_t) = q_t \Pr(D_{t+1} = \$60|s_t = bad) + (1 - q_t)\Pr(D_{t+1} = \$60|s_t = good)$.

The calculations above establish the Bayesian benchmark, which we use to incentivize subjects when they report their beliefs. We randomly pick one of the eight periods in which we elicit beliefs and WTP, and we then pay subjects based on either the beliefs question or the WTP question. If the beliefs question is randomly chosen, then we randomly select one of the possible dividend outcomes and pay subjects a \$3 bonus if their elicited probability estimate is within one percentage point of the objective probability of that outcome. For each percentage point that subjects deviate from the Bayesian prediction, we subtract 3 cents.

If instead the WTP question is randomly chosen, we implement a Becker-DeGroot-Marschak (BDM) mechanism, which is designed so that it is in the subject's best interest to accurately report their WTP. To implement the mechanism, we endow the subject with \$210 in experimental wealth, which can be used to purchase the right to next period's dividend. After the subject reports their WTP for next period's dividend, we draw a random price between \$60 and \$150. If the price that we draw is equal to or smaller than the WTP, the subject purchases the one period asset at the randomly drawn price. If the number is larger than the stated WTP, the subject does not purchase the asset. Subjects receive their remaining experimental wealth after any profits or losses from purchasing the asset. Each dollar in the experiment converts to \$0.01. Thus, subjects can receive a bonus of up to \$3 for the WTP question.

While it may be difficult for subjects to implement the Bayesian updating rule in (7), we emphasize that our main test in this experiment does not rely on subjects' ability to accurately compute q_t. Specifically, Prediction 1 states that if a subject has cognitive noise, then the elasticity between *subjective* expectations and WTP should be less than 1. Thus, our main test is independent of the expectation formation process. At the same time, the Bayesian benchmark is useful not only for eliciting incentive-compatible beliefs, but also to study any systematic differences between objective and subjective beliefs. The wedge between subjective and objective beliefs will turn out to be an an important predictor for the weak transmission across subjects.

We recruit 300 subjects from the online data collection platform, Prolific. The sample size and exclusion criteria are pre-registered on Aspredicted.org.[10] Subjects received \$2 for

[10]See https://aspredicted.org/6Z4_RLQ for the pre-registration document. After analyzing the data, the emphasis of our analysis changed to the weak transmission of beliefs to actions. We believe that including the initial pre-registration here is important for transparency, particularly about sample size and exclusion criteria. See Internet Appendix IA.3 for additional pre-registered analyses. Our analyses in Experiment 1 provide crucial motivation for the design of Experiments 2 and 3, which we also pre-register; details are provided in Sections III and IV.

completing the experiment, in addition to their bonus payment. The average completion time of the experiment was approximately 13 minutes, and the average earnings were \$4.39, including the \$2 participation fee.

B. Experimental results

B.1. Summary statistics

Our experiment with 300 subjects produces a panel dataset with 2,400 total observations (8 elicitations per subject). Table I provides summary statistics of the dataset where $\mathbb{E}^*$ denotes expectations under subjects' reported beliefs and $\mathbb{E}^b$ denotes the Bayesian expectation. Because all subjects face the same sequence of dividends, the time series of the Bayesian distribution is identical across subjects.

Table I
Summary Statistics from Experiment 1

		Mean	p25	p50	p75	SD	Min	Max
Subjective expected payoff	$\mathbb{E}^*[D]$	112.61	105.50	113.00	120.50	12.04	65.00	150.00
Deviation from Bayesian	$\mathbb{E}^*[D]/\mathbb{E}^b[D]$	1.01	0.96	1.02	1.08	0.10	0.59	1.32
Willingness to pay	P	95.15	80.00	95.70	110.00	20.88	60.00	150.00
Perceived volatility	$\text{Vol}^*[D]$	23.64	21.36	24.81	27.22	6.05	0.00	39.69
Bayesian volatility	$\text{Vol}^b[D]$	25.37	24.70	25.73	26.02	0.84	23.96	26.07

This table presents summary statistics for the main variables in our sample. The sample consists of 300 subjects and 8 elicitation periods, yielding 2,400 observations. $\mathbb{E}^*[D]$ is the subjective expected payoff, defined as the mean of a subject's reported dividend distribution. $\mathbb{E}^b[D]$ is the Bayesian expected payoff, defined as the mean of the Bayesian dividend distribution. P is the subject's reported willingness to pay for next period's dividend. Perceived volatility, $\text{Vol}^*[D]$, is the volatility of a subject's reported dividend distribution. Bayesian volatility, $\text{Vol}^b[D]$, is the volatility of the Bayesian distribution.

Table I reveals that subjects are quite accurate about the expected payoff on average, but some subjects report beliefs that strongly depart from the Bayesian benchmark. The average deviation from the Bayesian expectation, which we compute as $\mathbb{E}^*[D]/\mathbb{E}^b[D]$, is 1.01, where a value of 1.00 corresponds to no deviation from Bayesian expectations. The prices that subjects are willing to pay are, on average, lower than expected payoffs, which is consistent with risk aversion among our subjects. Because we elicit subjects' beliefs about the entire payoff distribution, we can measure perceived risk as the volatility of the elicited distribution. The average and median perceived volatility are fairly close to that of the Bayesian investor.

For all subsequent empirical analyses, we transform WTP and expectations to logs in order to be consistent with the conceptual framework in Section I. In particular, p denotes $\log P$, d denotes $\log \mathbb{E}^*[D]$, λ denotes $\text{Vol}^*[D]$, and r denotes $\log \mathbb{E}^*[R]$.

B.2. Testing for weak transmission

We begin by examining the strength of transmission from expectations to WTP. Recall that the frictionless benchmark outlined in Section I indicates that a 1-unit increase in d should translate into a 1-unit increase in p. We assume that the price of risk (γ) and the degree of weak transmission (x) are fixed over time within subjects – but can vary across subjects. In particular, all of our empirical results are based on mixed effects regressions with random slopes and a random intercept.

Column 1 of Table II reports that the sensitivity of WTP to payoff expectations is 0.634, which is significantly below one ($p < 0.001$). Figure 1 illustrates this result by showing that the best fitting line is much shallower than the 45-degree line. Column 2 indicates that the responsiveness of WTP to beliefs remains significantly below one after controlling for our assumed measure of risk, namely, conditional volatility.[11] The specification in Column 2 also confirms that subjects demand compensation for risk ($\gamma > 0$), as they are willing to pay less when risk is higher, holding the subjective expected payoff constant.

[11]In Appendix B, we show that under mild assumptions, our empirical test of weak transmission is robust to misspecifying the subject's perceived level of risk. Specifically, we show that omitting λ in a regression of p on d biases the estimate of x *upward* whenever (i) the price of risk is positive (e.g., subjects are risk averse) and (ii) the correlation between subjective payoff expectations d and perceived risk λ is negative.

Table II

Transmission of Expectations into Valuations

p	(1)	(2)
d	0.634***	0.610***
	(0.049)	(0.050)
λ		−0.195***
		(0.073)
Constant	1.539***	1.699***
	(0.233)	(0.242)
Observations	2,400	2,400

This table presents results from mixed effects regressions of log(WTP) (p) on log subjective expected payoff (d) and perceived volatility (λ). These regressions include a random effect for d and λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

We note that in this experiment, the cognitive default $\bar{p}_t$, is not observable. Thus, if the true model of decision-making is governed by Equation (4), then there is a potential omitted variable bias when regressing p_t on d_t and λ_t. In particular, our estimate of x will be biased if $\bar{p}_t$ and d_t are correlated. However, there are two reasons to believe that this potential omitted variable bias does not compromise our conclusion of a weak transmission. First, if $\bar{p}_t$ and d_t are correlated, they are likely to be positively correlated. When subjects report optimistic beliefs, they are likely to hold a high cognitive default. Conversely, when subjects report pessimistic beliefs, they are likely to hold a low cognitive default. A positive correlation between $\bar{p}_t$ and d_t will then induce an *upward* bias in our estimate of x, leading to a conservative estimate of weak transmission.

Second, later in the paper (Section IV) we report results from an additional experiment where we can directly observe $\bar{p}_t$. There, we find a zero correlation between $\bar{p}_t$ and d_t. Under the assumption that the zero correlation is stable across experiments, the omission of $\bar{p}_t$ from our regression in Table II will not affect our estimate of x.

Figure 1. Willingness to Pay and Subjective Expected Payoffs

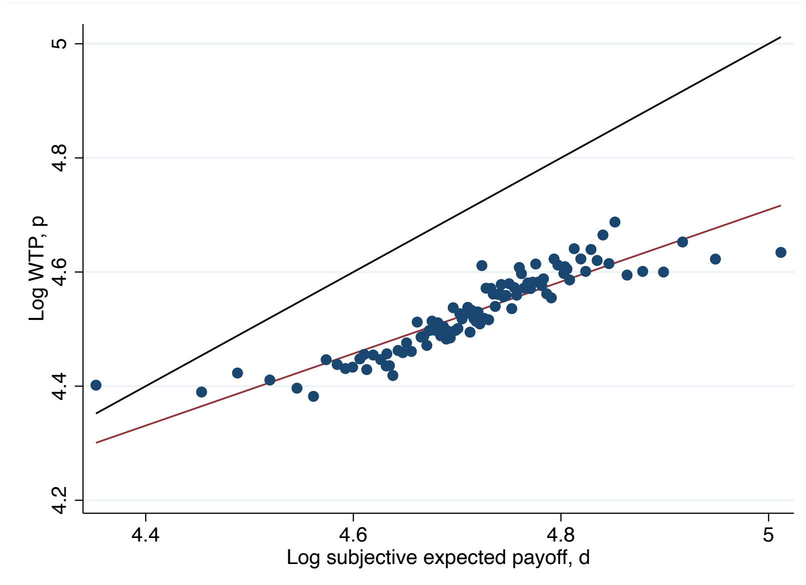

This figure is a binned scatter plot of log(WTP) (p) vs. log subjective expected payoff (d) controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300. The upper line is the 45-degree line.

B.3. Deviations from rational expectations and the link with weak transmission

When testing for weak transmission in the previous section, we do not impose any assumptions about whether subjective beliefs are rational or not. Because we have precise control over the stochastic process that generates dividends, here we can assess (i) the extent to which subjective beliefs depart from rational beliefs and (ii) whether beliefs that are closer to the rational benchmark are transmitted more vigorously into actions. Given the large literature on subjective expectation formation in finance (Adam and Nagel (2023)), it is useful to explore whether the severity of expectation errors has any bearing on the degree to which these errors are transmitted into investment actions.

For each subject we compute a measure of how well-calibrated their beliefs are using the absolute error summed across the eight elicitation periods. In particular, for each subject s, we compute: calibration $\text{error}_s = \sum_{t=1}^{8} |d_{st} - d_t^b|$ where d_s is the subjective expected payoff for subject s and d^b is the Bayesian expectation. The median calibration error across all 300 subjects is 0.552, and thus a substantial portion of our subjects report beliefs that have sizeable deviations from the Bayesian benchmark.

More interestingly, we can assess whether cross-sectional variation in beliefs tells us anything about cross-sectional variation in weak transmission. We define a dummy variable, $calib_{tot}$ that takes the value of 1 if a subject's calibration error is below the median – indicating their beliefs are relatively well-calibrated. We then re-estimate our basic regression of p on d, but allow for different slopes depending on whether the subject's beliefs are well-calibrated.

Column 1 of Table III shows that subjects with well-calibrated beliefs transmit their beliefs into actions much more strongly than subjects whose beliefs are not well-calibrated. The value of x for subjects with well-calibrated beliefs is 1.07, and we cannot reject the null that this coefficient is equal to the frictionless benchmark of 1 ($p = 0.363$). In contrast, the value of x for the sample of subjects whose beliefs are not well-calibrated is only 0.51, which is significantly below the frictionless benchmark of 1 ($p < 0.001$). To our knowledge, this is the first result demonstrating that the rationality of beliefs and the transmission of beliefs into actions are correlated in the cross-section.

Next, to better understand the structure of the expectation errors, for each subject we decompose the total calibration error into its fixed and variable components:

$$\text{calibration error}_s = \sum_{t=1}^{8} |d_{st} - d_t^b| = \underbrace{\left|\sum_{t=1}^{8} d_{st} - \sum_{t=1}^{8} d_t^b\right|}_{fixed} + \underbrace{\sum_{t=1}^{8} \left[|d_{st} - d_t^b| - 1/8 \left|\sum_{t=1}^{8} d_{st} - \sum_{t=1}^{8} d_t^b\right| \right]}_{variable}$$

The fixed component represents any optimism or pessimism bias across all eight elic-

Table III
Weak Transmission is Modulated by Expectation Errors

p	(1)	(2)
d	0.506***	0.393***
	(0.053)	(0.066)
d x $calib_{tot}$	0.562***	
	(0.092)	
d x $calib_{fix}$		0.339***
		(0.095)
d x $calib_{var}$		0.447***
		(0.098)
$calib_{tot}$	−2.661***	
	(0.434)	
$calib_{fix}$		−1.577***
		(0.445)
$calib_{var}$		−2.124***
		(0.462)
Constant	2.148***	2.672***
	(0.253)	(0.310)
Observations	2,400	2,400

This table presents results from mixed effects regressions of log(WTP) (p) on log subjective expected payoff (d) interacted with dummies that indicate below-median calibration errors for the total error ($calib_{tot}$), the fixed error ($calib_{fix}$), and the variable error ($calib_{var}$). Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

itations. The variable component captures overreaction or underreaction, and any other time-varying errors. The median calibration error for the fixed component is 0.226, whereas the median calibration error for the variable component is 0.256. We find that across subjects, the fixed and variable components are significantly negatively correlated at −37%. These results suggest that expectation errors are not exclusively driven by any one theory of non-rational expectations, such as overreaction or pessimism. Instead, both the fixed and variable components systematically contribute to expectation errors.[12]

To understand whether one source is predominantly responsible for the empirical link with weak transmission, we re-estimate the regression from Column 1 of Table III, but we use two different dummy variables: $calib_{fix}$ and $calib_{var}$. The dummy variable $calib_{fix}$ takes a value of 1 if the subject is below the median in their fixed calibration error. Similarly, the dummy variable $calib_{var}$ takes a value of 1 if the subject is below the median in their variable calibration error. Column 2 shows that *both* components independently and significantly explain variation in weak transmission.

To summarize, we find that there is substantial variation across subjects in the degree to which their subjective expectations coincide with rational expectations. This variation is empirically connected to the observed weak transmission: those subjects who report beliefs closer to the Bayesian benchmark are the same subjects who transmit their beliefs more strongly into actions. Moreover, there appear to be at least two sources of deviations from rational expectations, which we capture by a fixed and variable component of expectation errors. Both components explain a portion of the variation in weak transmission. In the next section, we shut down the possibility that subjects hold non-rational expectations, and assess whether this restriction causally increases the transmission of beliefs into actions.

[12]Because we do not elicit a term structure of beliefs, we cannot make strong statements about the precise type of expectational error captured by the variable component. For example, in order to test for overreaction or underreaction using the regression framework of Coibion and Gorodnichenko (2015), we would need data on how beliefs evolve for a stochastic dividend at fixed maturity. Since the main goal in this section is to test for weak transmission *after* expectations are formed, we elicit only subjects' 1-period ahead beliefs and the associated valuation.

III. Experiment 2: Exogenous manipulation of cognitive noise

In this section, our goal is to measure the causal effect of cognitive noise on the transmission of beliefs into actions. Our approach to manipulating cognitive noise is to experimentally reduce the complexity of the decision environment. This method of manipulating cognitive noise builds on recent work by Enke and Graeber (2023), who show that reducing the complexity of a decision problem leads to lower self-reported measures of cognitive noise. We argue that a decision environment in which subjects must learn from past data about the distribution of future dividends is more complex than one in which subjects are endowed with objective beliefs.

By comparing behavior in an environment where subjects are endowed with beliefs with behavior from our first experiment, we can assess whether cognitive noise causally affects transmission strength. Moreover, by testing whether the transmission strength remains below the frictionless benchmark, we can assess whether other sources besides belief uncertainty are important drivers of cognitive noise.

A. Experimental design

In Experiment 1, subjects face a learning problem in which realized dividends can be used to form Bayesian beliefs about the next period's dividend. While subjects are endowed with all information about the data generating process, implementing the Bayesian updating rule is complex. Thus, subjects may be cognitively uncertain about their own beliefs, and we hypothesize that this cognitive uncertainty dampens the transmission of beliefs to WTP.

In order to test this hypothesis, here we conduct an experiment that is identical to Experiment 1, except we endow subjects with objective beliefs about the next period's dividend. Subjects do not need to learn because we explicitly provide them with the objective payoff distribution. Our manipulation is meant to reduce cognitive noise which, in turn, should increase the strength of the transmission of beliefs to WTP.

Perhaps the most natural experimental design would involve simply endowing subjects with the Bayesian beliefs from Experiment 1. An issue however, is that the WTP elicited in such a design would be based on beliefs that differ from the subjective beliefs reported by subjects in Experiment 1. Any difference in behavior could thus be due to differences in beliefs, rather than a difference in the objectivity of those beliefs. Thus, we would not be able to identify cognitive noise as a channel through which WTP becomes more responsive to beliefs.

To sidestep this concern, we design an experiment in which we recruit a new set of 300 subjects, and each subject is uniquely matched to a subject from Experiment 1. The new subject in Experiment 2 inherits the beliefs reported by her matched partner. That is, the subjective beliefs reported by the subject in Experiment 1 become the objective beliefs for the subject in Experiment 2. Subjects in Experiment 2 are not told anything about the source of such beliefs, or even about the existence of Experiment 1. Instead, we incentivize subjects from Experiment 2 to report their WTP for an asset that pays a dividend according to the objective distribution that we present them. We provide screenshots of this experiment in Internet Appendix IA.2.

This design allow us to test Prediction 2: reducing cognitive noise will strengthen the transmission of beliefs to WTP. As in our previous experiment, here subjects are incentivized using the BDM mechanism, and we randomly select one of the questions for payment at the end of the experiment. Note that subjects in Experiment 2 therefore answer only 8 questions (compared with 16 in Experiment 1). To keep the incentives per question similar across experiments, we cut the bonus incentive in Experiment 2 in half compared with Experiment 1. This is important because larger incentives could lead to lower cognitive noise in Experiment 2.

As in Experiment 1, we recruit subjects from Prolific and pre-register the experiment on Aspredicted.org.[13] Subjects received \$2 for completing the experiment, in addition to their bonus payment. The average completion time of the experiment was 6 minutes, and the average earnings were \$3.06 including the \$2 participation fee.

B. *Experimental results*

We begin by testing whether our experimental manipulation strengthens the transmission of beliefs to WTP. Column 1 in Table IV presents results from regressions of WTP on payoff expectation using data from both experiments, where "Exp2" is a dummy variable that equals one if and only if the observation is from Experiment 2. Consistent with Prediction 2, the transmission of beliefs to WTP is causally strengthened when we endow subjects with objective beliefs, thereby eliminating the need for subjects to form their own beliefs based on past data. Among subjects in Experiment 2, the coefficient on d is 0.877 (= 0.634 + 0.243), which is significantly higher than 0.634 in Experiment 1.

The stronger transmission of payoff expectations to WTP remains significant after controlling for risk as shown in Column 2 of Table IV. The coefficient on d significantly increases from 0.610 in Experiment 1 to 0.817 in Experiment 2.[14] These results suggest that cognitive noise from expectation formation explains about half of the underreaction of WTP to payoff expectations.

We also find that decreasing cognitive noise increases the reaction of WTP to risk perception. Column 2 of Table IV shows that the negative loading on risk increases in magnitude when going from Experiment 1 to 2. The coefficient on risk more than doubles in magnitude,

[13] For pre-registration details, see: https://aspredicted.org/NWL_YML

[14] As discussed in Section II, there is a potential omitted variable bias because $\bar{p}_t$ is not observable in this experiment. However, under the assumption that the correlation between $\bar{p}_t$ and d_t is the same across Experiments 1 and 2, there will be no bias in estimating the difference in x across Experiments 1 and 2.

Table IV
Transmission Strength Increases When Beliefs are Endowed

p	(1)	(2)
d	0.634***	0.610***
	(0.049)	(0.050)
d x Exp2	0.243***	0.207***
	(0.064)	(0.066)
λ		−0.188***
		(0.073)
λ x Exp2		−0.221*
		(0.121)
Exp2	−1.134***	−0.910***
	(0.302)	(0.322)
Constant	1.539***	1.695***
	(0.233)	(0.242)
Observations	4,800	4,800

This table presents results from mixed effects regressions of log(WTP) (p) on log subjective expected payoff (d) and perceived volatility (λ), combining data from Experiments 1 and 2. These regressions include a random effect for d and λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

though the effect is only statistically significant at the 10% level (the large standard error here is due in part to the strong negative correlation between λ and d). This result is consistent with our conceptual framework, as Equation (4) shows that the impact of perceived risk on WTP is also dampened by a factor of x.

While Table IV reveals that the transmission of expectations to WTP is significantly stronger in Experiment 2, we emphasize that the transmission strength is still not as strong as predicted by the frictionless benchmark. In particular, the estimated value of x of 0.817 in Experiment 2 is still significantly below the frictionless benchmark of 1 ($p < 0.001$). Thus, other features of the decision-making process besides belief formation – such as the integration of beliefs with preferences – are likely to be implemented with cognitive noise, which in turn sustains the weak transmission of expectations into actions.

To summarize the main finding from Experiment 2, we provide causal evidence that WTP is substantially more responsive to beliefs when subjects price an asset based on objective – rather than subjective – beliefs. Our results suggest that subjects rely much less on a cognitive default parameter when they are endowed with objective beliefs. In our final experiment, we develop a method to measure the cognitive default parameter, and we subsequently provide a novel test of the cognitive noise mechanism.

IV. Experiment 3: Measuring the cognitive default

In this section, we report results from a third experiment that employs a novel method for eliciting the cognitive default. Armed with direct measures of the cognitive default at the subject-period level, we can provide an additional and sharp test of theory.

The logic of the main empirical test we run in this section is as follows. According to Equation (4), when the subject is forming her valuation of the asset, she applies a weight of $(1 - x)$ to the cognitive default $\bar{p}_t$, and she applies a weight of x to d_t. In the previous two sections, we were concerned with testing whether the estimated x is less than 1, and we assumed that the subject must be putting weight $(1 - x)$ on the cognitive default. Here, we create a setting in which we can observe both d_t and $\bar{p}_t$, which enables us to test whether the weights on these two variables are both positive and sum to 1, as predicted by theory.

A. Experimental design

The design of this experiment is nearly identical to that in Experiment 2. The only difference is that we add one additional question in each of the eight periods. Specifically, in each period, *before* the distribution of dividends is revealed to the subject, we ask the subject to report their willingness to pay for the asset.[15] This reported valuation represents the subject's cognitive default: it is the price she expects to pay, before drawing her noisy cognitive signal about the specific problem she faces. Thus, for each subject and each period, we collect a pair of valuations: $\{\bar{p}_t, p_t\}$.[16]

We elicit both the cognitive default and the WTP in an incentive-compatible fashion. We instruct subjects that we will randomly select one of the sixteen valuations they report (eight WTPs and eight cognitive defaults). As in the previous two experiments, we then implement a BDM mechanism given the subject's reported valuation and the realized dividend from the asset. Note that the BDM is plausible to implement, even if the distribution of dividends is not yet displayed to subjects. This allows us to use the same BDM procedure regardless of whether we randomly choose a cognitive default or a WTP for payment at the end of the experiment.

We recruit a new set of 300 subjects from Prolific for Experiment 3. Each subject is then uniquely matched to a subject from Experiment 2, and she sees the same sequence of eight dividend distributions as her matched partner from Experiment 2. This feature of the design enables us to test whether the *act of eliciting* the cognitive default affects WTP, since the beliefs that are used to price the assets are identical across the two experiments.

We pre-registered the experiment on Aspredicted.org.[17] Subjects received $2 for completing the experiment, in addition to their bonus payment. Since Experiment 3 contains sixteen questions, we doubled the size of the bonus compared to Experiment 2 (which contains only

[15] We provide screenshots of this experiment in Internet Appendix IA.2.

[16] Note that we designed this third experiment by building on the simpler design of Experiment 2 – where we still obtain weak transmission – rather than the more complex design of Experiment 1. We reasoned that interpreting the cognitive default would be cleaner in an environment where beliefs are endowed to the subject and where the elicitation of the cognitive default can be administered immediately before revealing the objective payoff distribution to the subject.

[17] For pre-registration details, see: https://aspredicted.org/MYN_NMS.

eight questions). This implies that the expected bonus per question should be similar across Experiments 1, 2, and 3. The average completion time of the experiment was approximately 9 minutes, and the average earnings were \$4.12 including the \$2 participation fee.

B. *Experimental results*

Our goal in this section is to test whether the measured cognitive default correlates with valuation in the manner predicted by theory. We estimate the regression model specified in Equation (4), and we conduct two pre-registered empirical tests: (i) does the measured cognitive default positively correlate with valuation? and (ii) do the estimated coefficients on $\bar{p}_t$ and d_t sum to 1? Column 1 of Table V provides results from a regression of valuation on expected payoff and the cognitive default. We find that the coefficient on $\bar{p}$, which denotes the log of the elicited cognitive default, is 0.220 and is significantly above 0 ($p < 0.001$). This indicates that the willingness to pay that subjects report before seeing the dividend distribution is a significant predictor of valuation *even after* the dividend distribution is displayed. This is consistent with our main hypothesis, as cognitive noise corrupts the subject's ability to precisely report her valuation after the dividend distribution is displayed, and thus she optimally shades her reported valuation towards the cognitive default.

We provide a second test of the mechanism by testing whether the coefficients on $\bar{p}$ and d sum to 1. Table V shows that the sum of the coefficients on $\bar{p}$ and d is 1.07; we cannot reject the null hypothesis that the sum is equal to 1 ($p = 0.19$). This result provides us with greater confidence in interpreting the coefficient on d as a ratio of variances. Recall that in our model, $x = \sigma_p^2/(\sigma_p^2 + \sigma_\epsilon^2)$, and thus the coefficient on d represents the amount of noise in the cognitive signal, relative to the total noise from the prior and the cognitive signal.

One potential concern with these results is that the elicitation of the cognitive default may causally affect WTP. In particular, it is plausible that eliciting the cognitive default immediately before asking subjects for WTP may artificially inflate the weight that subjects

Table V
Valuation Depends on Cognitive Default

p	(1)	(2)
d	0.847***	0.829***
	(0.045)	(0.045)
$\bar{p}$	0.220***	
	(0.030)	
d x Exp2		−0.010
		(0.063)
λ	−0.372***	−0.368***
	(0.097)	(0.099)
λ x Exp2		−0.033
		(0.138)
Exp2		0.033
		(0.304)
Constant	−0.316	0.747***
	(0.242)	(0.217)
Observations	2,400	4,800

This table presents results from mixed effects regressions. Column 1 regresses log(WTP) (p) on log subjective expected payoff (d), the log cognitive default ($\bar{p}$), and perceived volatility (λ) using data only from Experiment 3. Column 2 combines data from Experiments 2 and 3 and regresses log(WTP) (p) on log subjective expected payoff (d) and perceived volatility (λ). All regressions include a random effect for d and λ, as well as for the intercept. Column 1 additionally includes a random effect for $\bar{p}$. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

attach to the cognitive default when reporting valuations. Because the only difference across Experiments 2 and 3 is the elicitation of the cognitive default, we can compare WTP across experiments to gain insight on the causal effect of eliciting the cognitive default. If eliciting the cognitive default caused subjects to inflate its weight when forming valuations, then this should lead to a *lower* weight on d in Experiment 3. This, in turn, should be reflected in a positive coefficient on the interaction term, d x Exp2. Instead, we find that the estimated coefficient on the interaction is not significantly different from 0. Similarly, Column 2 shows that subjects rely on volatility to the same extent in both Experiment 2 and Experiment

3. In sum, the results in Column 2 provide reassurance that the elicitation of the cognitive default is not artificially changing the valuation process.[18]

Given that the cognitive default does explain significant variation in reported valuations, it is useful to explore which factors drive variation in the cognitive default. To start, we note that the bulk of the variation in the cognitive default is explained by a subject fixed effect. Specifically, we find that 77% of the variation of the cognitive default can be explained by a subject level fixed effect (see Internet Appendix IA.4). However, the results in Table V are identified off of within-subject variation, which suggests that time-varying observables may explain additional variation in reported cognitive defaults.

To explore this direction, we regress the cognitive default on lagged d, lagged p, and the current level of d. The results are shown in Table VI. We see that the lagged WTPs - but not lagged d - significantly explain variation in the cognitive default. At the same time, including these lagged variables in an OLS regression barely increases the explained variation in the cognitive default. When analyzing elicitations 3 - 8 (the same sample used in Table VI), the R^2 from an OLS regression with only subject fixed effects is already at 85.4%, and this increases to 85.8% when we add d_{t-1}, d_{t-2}, p_{t-1}, and p_{t-2} to the regression (see Columns 2 and 3 of Table IA.2).

One way to interpret the above results is that the cognitive default reflects a wide variety

[18] As an alternative approach to understanding the impact of eliciting the cognitive default on subsequent valuation, we pre-registered an exploratory analysis that examines whether the time series of cognitive defaults from Experiment 3 can predict valuations from Experiment 2. To implement the test, we re-estimate the regression from Column (1) of Table V, but replace the WTPs with those reported by subjects from Experiment 2. When restricting to periods 5-8, as outlined in our pre-registration, we find that the estimated coefficient on $\bar{p}$ is 0.007 and is not significantly different from zero ($p = 0.84$). The fact that cognitive defaults from Experiment 3 do not predict valuations in Experiment 2 is consistent with our finding, described in more detail below, that most of the variation in the cognitive default is subject-specific and does not depend on lagged d (which is common across Experiments 2 and 3).

Table VI

Explaining Variation in the Cognitive Default

$\bar{p}_t$	(1)	(2)	(3)
d_t	0.019	0.024	0.028
	(0.028)	(0.028)	(0.028)
d_{t-1}		−0.035	−0.031
		(0.033)	(0.032)
d_{t-2}			−0.005
			(0.027)
p_{t-1}		0.144***	0.144***
		(0.027)	(0.026)
p_{t-2}			0.090***
			(0.020)
Constant	4.379***	3.860***	3.430***
	(0.130)	(0.213)	(0.258)
Observations	1,800	1,800	1,800

This table presents results from mixed effects regressions of the log cognitive default ($\bar{p}$) on current period and lagged log subjective expected payoff (d) and lagged log(WTP)(p). The sample includes elicitation periods 3-8. These regressions include a random effect for d and p, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

of personal-level characteristics that are stable within our experiment: risk aversion, past experience with investing, and any exposure to previous experiments. These stable features explain nearly all of the variation in observed cognitive defaults. While the time-variation in cognitive defaults is minimal, our result that lagged valuations significantly predict the current cognitive default may prove useful in future research that seeks to provide a theory of how the cognitive default evolves over time. We speculate that properties of human memory will play an important role in understanding which features contribute to the time-varying cognitive default (Bordalo, Gennaioli, and Shleifer (2020)).

V. Discussion

In this section, we discuss the implications of our results for asset pricing models and draw connections with the broader literature on survey data. We also discuss in more detail the psychological source of weak transmission in our experiments, and we discuss limitations of our experimental approach.

A. Implications for asset pricing models

A.1. Incorporating weak transmission into investor decisions

Over the past decade, there has been an enormous amount of theoretical and empirical work on expectation formation in finance (Adam and Nagel (2023)). Several researchers have proposed models with the goal of jointly explaining asset prices and survey expectations (e.g., Barberis et al. (2015), Hirshleifer, Li, and Yu (2015), Jin and Sui (2022)). While these models formalize the subjective expectation formation process in a psychologically grounded manner, they retain the standard assumption that investors fully act on their subjective beliefs.

To better place our contribution in this literature, it is useful to decompose the investor's decision process into two stages: (i) expectation formation and (ii) action selection. While a bulk of the literature has studied the first stage, we take expectations as given and focus on how these expectations propagate into actions. One key finding that emerges from all three of our experiments is that the transmission of beliefs to actions is far below 1-to-1. We believe that this result should motivate future theoretical work that explicitly incorporates weak transmission into the investor's decision process.[19]

One concern, however, is that injecting deviations from rationality into both expectation

[19]For related work in macroeconomics, see Khaw, Stevens, and Woodford (2017) for a theoretical model that incorporates inattentive adjustment of actions.

formation and expectation transmission gives the modeler too much flexibility. Indeed, the number of non-rational expectations models is already large, and adding an extra degree of freedom does not help the case for parsimony. However, there is an intriguing possibility that departures from rational expectations are connected to how these expectations propagate into actions. In Table III, we show that subjects who state beliefs closer to the rational benchmark are the same subjects whose valuations are more sensitive to these beliefs. Furthermore, Andries et al. (2022) show that experimental subjects transmit beliefs into actions less vigorously when they hold (overly) extrapolative expectations compared to rational expectations. While more empirical data is needed, together, these results suggest that the belief formation process can partially constrain the degree of weak transmission.

Our results also provide guidance for asset pricing models that feature learning and state uncertainty. In Experiment 1, we implement an imperfect information environment in which subjects receive noisy signals about the state in the form of realized dividends. Using these signals, subjects are incentivized to form subjective beliefs about the conditional distribution of payoffs. We find that subjects who report beliefs that are closer to the rational benchmark also transmit these beliefs more vigorously into actions. This cross-sectional result may provide useful insights for modeling the actions of investors depending on the rationality of their belief formation process. For example, in models where investors learn through Bayesian inference (e.g., Veronesi (1999), Johannes, Lochstoer, and Mou (2016), Ghaderi, Kilic, and Seo (2022)), our results suggest that beliefs should be strongly transmitted into actions. On the other hand, in models where learning departs from Bayesian inference or investors hold perpetually misspecified models (e.g., Jin and Sui (2022), Nagel and Xu (2022)), our results suggest that investor beliefs will be transmitted weakly into actions.[20]

[20]In Internet Appendix IA.5 we provide further results on the implication of weak transmission for basic asset pricing tests.

A.2. Cognitive noise as a source of inelastic demand

The data we produce connects to a recent literature which argues that inelastic demand can explain important features of the aggregate stock market such as excess volatility (Koijen and Yogo (2019), Gabaix and Koijen (2022)). While there are a variety of institutional reasons why investors may exhibit inelastic demand, cognitive noise may provide one psychological source of the inelasticity.

Here we build on our conceptual framework from Section I and sketch a simple alternative specification that imposes cognitive noise over asset demand, rather than asset valuation. We begin with a baseline case where there is a single asset in fixed supply. Following Haddad, Huebner, and Loualiche (2022), we define the investor's log demand as

$$\begin{aligned} q &= \bar{q} - \zeta(p - p_e) \\ &= \bar{q} - \zeta p_{\Delta}, \end{aligned} \tag{8}$$

where p is the log asset price, p_e is a log baseline price, ζ is the elasticity of demand, and p_{Δ} denotes the deviation between the log asset price and the log baseline price. The variable $\bar{q}$ captures the component of demand that is not driven by deviations between the market price and the baseline price (e.g., the risk profile of the asset). When ζ is very high, investors will absorb demand shocks by aggressively trading against p_{Δ}. Conversely, when ζ is very low, idiosyncratic demand shocks will impact market prices as investors will not aggressively trade against p_{Δ}.

Suppose now that the investor does not have access to her true demand q, but only has access to a noisy cognitive signal of this demand: $q^0 = q + \omega$, where ω is distributed according to $N(0, \sigma_{\omega}^2)$. The cognitive noise term, ω, is meant to capture the difficulty that the investor faces in computing her precise asset demand. This difficulty could stem from uncertainty about how far the market price is from the baseline price. The investor may also be uncertain about her own optimal value of ζ, given her beliefs about the demand elasticity of other market participants (Haddad, Huebner, and Loualiche (2022)).[21] Finally, even if the investor has precise access to the demand inputs, $\bar{q}$, ζ, and p_Δ, she may find it difficult to translate these inputs into an exact quantity of shares demanded.

Similar to the setup in Section I, we assume the investor has prior beliefs about her true asset demand. We further impose that these beliefs are centered around $\bar{q}$, and are given by $N(\bar{q}, \sigma_q^2)$. Under this scenario, the investor's demand is given by:

$$\begin{aligned} q &= (1 - x_q)\bar{q} + x_q q^0 \\ &= (1 - x_q)\bar{q} + x_q(\bar{q} - \zeta p_\Delta + \omega) \\ &= \bar{q} - x_q \zeta p_\Delta + x_q \omega, \end{aligned} \tag{9}$$

where $x_q = \sigma_q^2/(\sigma_q^2 + \sigma_\omega^2)$. Because $0 < x_q < 1$, Equation (9) implies that demand is compressed towards $\bar{q}$, relative to the noiseless demand in Equation (8). Moreover, as the investor faces more cognitive noise (i.e., as σ_ω^2 increases), her demand becomes more inelastic.

Of course, a key assumption in the above argument is that the cognitive default coincides with $\bar{q}$. One interpretation of this assumption is that the investor believes that p_Δ is closer to 0 than its true value. For example, if p_e represents the fundamental value of the asset, then the assumption may be further interpreted as the investor having an overly strong belief in market efficiency. A natural step for future research would be to measure the cognitive default, as we do in Experiment 3, but in a situation where subjects report demand, rather

[21] In our experiments, subjects are price takers and their payoffs do not depend on the behavior of other subjects in the experiment. In a more complex setting where investors need to form beliefs about the elasticity of other traders, we hypothesize that cognitive noise would be even more pronounced. Some evidence for this conjecture comes from Frydman and Nunnari (2023), who show experimentally that cognitive noise is larger when a subject plays a strategic game against another human subject, compared to when she plays the same game against a computerized opponent.

than valuation.

B. Sources of the weak transmission

In the field, there are several potential reasons for the weak transmission of beliefs to actions. For example, Giglio et al. (2021a) discuss frictions such as capital gains taxes, institutional settings of retirement plans, and infrequent trading. Our experiment rules out such institutional frictions by design, and allows us to identify the weak transmission as driven by a psychological friction. In Experiment 1, where subjects need to form subjective beliefs, it is likely that a portion of cognitive noise arises from uncertainty about expectations, perhaps because subjects have difficulty implementing Bayes' rule (Kuhnen (2015)), Ben-David, Fermand, Kuhnen, and Li (2023)). But importantly, in Experiments 2 and 3, we show that shutting down uncertainty about beliefs still leads to weak transmission. We speculate that the cognitive noise in Experiments 2 and 3 arises primarily from integrating beliefs about payoffs with perceived risk to arrive at a valuation.

Another recently proposed source of weak transmission comes from Barberis and Jin (2023). Those authors develop a model of investor behavior based on reinforcement learning and argue that it can explain a variety of facts about financial markets, including the disconnect between beliefs and portfolios. While the psychology is quite different across models, cognitive noise and reinforcement learning both provide a foundation for the belief-action disconnect that arises from the investor's decision process, rather than from institutional constraints.

C. Limitations

The one period nature of the asset in our experiments is useful because it allows us to see how valuation relates to expectations in a simple setting. Indeed, we find clear evidence of a weak transmission of beliefs to actions, even when there is no need for subjects to form expectations over long horizons. Yet this simplicity also means that our analyses cannot speak directly to other previously documented facts about subjective expectations from the field.

For example, one of the most salient facts from the survey literature is that investors extrapolate recent returns when forming expectations about future returns (Greenwood and Shleifer (2014), Barberis et al. (2015)). One reason we do not analyze this dimension of the data in our experiments is because the degree of extrapolation, and more generally, expectational errors, may depend on the horizon of the forecast (Giglio and Kelly (2018), Da, Huang, and Jin (2021), De Silva and Thesmar (2023)). One opportunity for future work is to enrich the experimental design we present here by having subjects price an asset that delivers a long stream of cash flows – rather than a one period dividend strip. For example, one could integrate into our design the experimental method from Afrouzi et al. (2023), which elicits expectations along the term structure. This would further enable testing of other important phenomena, including the dividend-price ratio and its ability to predict returns of long-duration assets such as aggregate equity.

Finally, our experiments are designed to probe the mechanism that generates the weak transmission of beliefs to actions that has been observed in the field. The evidence we have presented points to cognitive noise as an important driver of weak transmission in the lab. As with most laboratory studies, we adopt the assumption that the mechanism we study in a controlled environment is similar to the one that partially drives behavior in the field. Of course, without further tests, we cannot rule out the possibility that the weak transmission observed in the field may be driven by alternative mechanisms. On the other hand, because investors arguably face more complex problems in the field, it is also possible that the role of cognitive noise for investor decision-making is even larger outside the lab.

VI. Conclusion

Survey data on subjective beliefs have recently opened up a vibrant area of research in asset pricing (Adam and Nagel (2023)). Subjective beliefs data offer the promise of disciplining models using the expectations that investors actually report, rather than the rational expectations that investors are typically assumed to hold. Our paper contributes to this agenda by experimentally studying *why* investors may not fully act on their reported beliefs. Our results suggest that researchers should approach survey data with some caution, if the intention is to use subjective expectations data to understand quantitative patterns in asset prices.

At the same time, a better understanding of the mechanism that generates weak transmission can aid in interpreting the implications of subjective beliefs data. Experiments 2 and 3 were designed precisely to provide such tests of the mechanism. Our results reveal that a long-standing idea from psychology – cognitive noise – is at least partially responsible for weak transmission. We interpret the cognitive noise in our experiments as arising from a combination of uncertainty about expectations, uncertainty about perception of risk, and the cognitive process of integrating these quantities to arrive at an asset's valuation.

One potential direction for future research is to incorporate a weak transmission mechanism into belief-based asset pricing models. Under this approach, subjective expectations data can still be harnessed, but the quantitative implications will be constrained by the degree to which these expectations are actually incorporated into asset demand. Another path forward is to better understand the relative importance that institutional frictions play in generating weak transmission in the field. Because institutional frictions are shut down in our experiments by design, it is plausible that the passthrough from beliefs to actions may be even weaker outside the laboratory. Thus, any asset pricing implications that stem from weak transmission may be even more pronounced when analyzing data from the field.

APPENDIX

A. Gaussian signal extraction

We adapt the basic Bayesian signal extraction studied by Gabaix (2019) to our conceptual framework. Suppose that the agent's objective is to minimize the squared distance between her true willingness to pay p^* and the willingness to pay p conditional on her noisy signal $p^0 = p^* + \epsilon$, where ϵ is normally distributed with mean 0 and variance σ_ϵ^2:

$$\max_p \mathbb{E}[-1/2(p-p^*)^2|p^o]. \tag{A.1}$$

Hence the optimality condition is $\mathbb{E}[p-p^*|p^o]=0$. Because ϵ has a zero mean, the prediction about p^* conditional on the signal p^0 is $\mathbb{E}[p^*|p^o]=(1-x)\bar{p}+xp^o$ where the dampening factor is given by:

$$x = \frac{\sigma_p^2}{\sigma_p^2+\sigma_\epsilon^2}. \tag{A.2}$$

As the variance of the noisy signal increases, the agent optimally puts more weight on the default $\bar{p}$.

B. Estimating x from willingness to pay and payoff expectations

In the following, we show that the univariate relation between p and d results in a upward-biased estimate of x if payoff expectation d and perceived risk λ are negatively correlated. Suppose that the relation between d and λ is affine and given by $d_t = \alpha + \beta\lambda_t + \eta_t$ where α and β are constants, and η_t represents variation in d_t that is orthogonal to λ_t. Rearranging results in

$$\lambda_t = -\frac{\alpha}{\beta} + \frac{1}{\beta}d_t - \frac{1}{\beta}\eta_t, \tag{A.3}$$

which can be plugged into (4) to obtain

$$p_t = (1-x)\bar{p}_t + \frac{x\gamma\alpha}{\beta} + x\left[1 - \frac{\gamma}{\beta}\right] d_t + \frac{x\gamma}{\beta}\eta_t. \tag{A.4}$$

As a result, the coefficient of p on d is $x\left[1-\frac{\gamma}{\beta}\right]$ which is larger than x if $\gamma > 0$ and $\beta < 0$. The following figure shows that $\beta < 0$ in the subjective beliefs data.

Figure A.1. Subjective Expected Payoffs and Perceived Risk

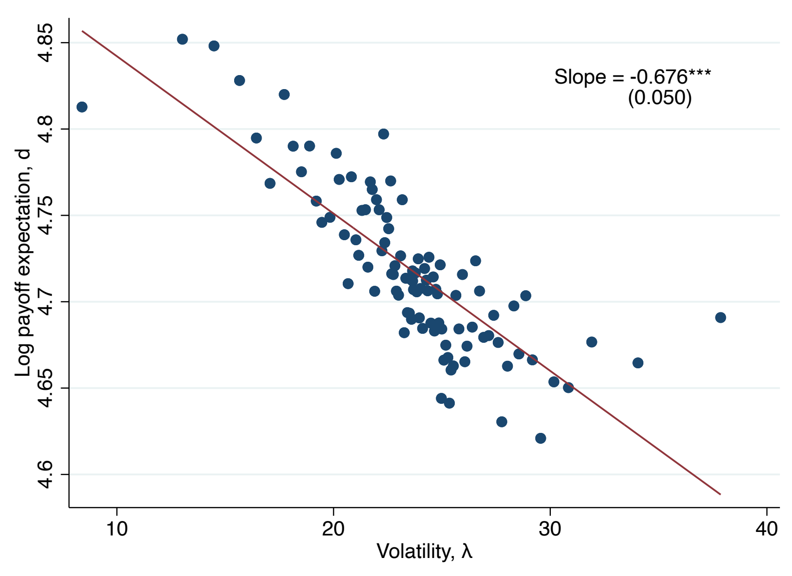

This figure is a binned scatter plot of log subjective expected payoff (d) and perceived volatility (λ) controlling for subject fixed effects. The reported slope results from a mixed effects regression of d on λ. The regression includes a random effect for λ as well as for the intercept. The displayed coefficient and standard error for λ are multiplied by 100. The standard error in parentheses is clustered at the subject level. The sample size is 2,400 and the number of subjects is 300. The data are from Experiment 1.

REFERENCES

Adam, Klaus, and Stefan Nagel, 2023, Expectations data in asset pricing, in *Handbook of Economic Expectations*, 477–506 (Elsevier).

Afrouzi, Hassan, Spencer Yongwook Kwon, Augustin Landier, Yueran Ma, and David Thesmar, 2023, Overreaction in expectations: Evidence and theory, *Quarterly Journal of Economics* forthcoming.

Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D Shapiro, 2020, Heterogeneity in expectations, risk tolerance, and household stock shares: The attenuation puzzle, *Journal of Business & Economic Statistics* 38, 633–646.

Amromin, Gene, and Steven A Sharpe, 2014, From the horse's mouth: Economic conditions and investor expectations of risk and return, *Management Science* 60, 845–866.

Andries, Marianne, Milo Bianchi, Karen K Huynh, and Sébastien Pouget, 2022, Return predictability, expectations, and investment: Experimental evidence, SSRN Working paper.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1–24.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2018, Extrapolation and bubbles, *Journal of Financial Economics* 129, 203–227.

Barberis, Nicholas, and Lawrence Jin, 2023, Model-free and model-based learning as joint drivers of investor behavior, Working paper.

Ben-David, Itzhak, Elyas Fermand, Camelia M Kuhnen, and Geng Li, 2023, Extrapolative uncertainty and household economic behavior, NBER Working paper.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019, Diagnostic expectations and stock returns, *Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *Quarterly Journal of Economics* 135, 1399–1442.

Brunnermeier, Markus, Emmanuel Farhi, Ralph SJ Koijen, Arvind Krishnamurthy, Sydney C Ludvigson, Hanno Lustig, Stefan Nagel, and Monika Piazzesi, 2021, Perspectives on the future of asset pricing, *Review of Financial Studies* 34, 2126–2160.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.

Da, Zhi, Xing Huang, and Lawrence J Jin, 2021, Extrapolative beliefs in the cross-section: What can we learn from the crowds?, *Journal of Financial Economics* 140, 175–196.

De Silva, Tim, and David Thesmar, 2023, Noise in expectations: Evidence from analyst forecasts, NBER Working paper.

Dehaene, Stanislas, 2011, *The Number Sense* (Oxford University Press, Oxford, United Kingdom).

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

Enke, Benjamin, and Thomas Graeber, 2023, Cognitive uncertainty, *Quarterly Journal of Economics* forthcoming.

Findling, Charles, and Valentin Wyart, 2021, Computation noise in human learning and decision-making: origin, impact, function, *Current Opinion in Behavioral Sciences* 38, 124–132.

Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel, 2014, Using neural data to test a theory of investor behavior: An application to realization utility, *Journal of Finance* 69, 907–946.

Frydman, Cary, and Lawrence Jin, 2022, Efficient coding and risky choice, *Quarterly Journal of Economics* 137, 161–213.

Frydman, Cary, and Salvatore Nunnari, 2023, Coordination with cognitive noise, SSRN Working paper.

Gabaix, Xavier, 2019, Behavioral inattention, in *Handbook of Behavioral Economics: Applications and Foundations 1*, 261–343 (Elsevier).

Gabaix, Xavier, and Ralph SJ Koijen, 2022, In search of the origins of financial fluctuations: The inelastic markets hypothesis, NBER Working Paper.

Ghaderi, Mohammad, Mete Kilic, and Sang Byung Seo, 2022, Learning, slowly unfolding disasters, and asset prices, *Journal of Financial Economics* 143, 527–549.

Giglio, Stefano, and Bryan Kelly, 2018, Excess volatility: Beyond discount rates, *Quarterly Journal of Economics* 133, 71–127.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021a, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021b, The joint dynamics of investor beliefs and trading during the COVID-19 crash, *Proceedings of the National Academy of Sciences* 118.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714–746.

Haddad, Valentin, Paul Huebner, and Erik Loualiche, 2022, How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing, SSRN Working paper.

Hartzmark, Samuel M, Samuel D Hirshman, and Alex Imas, 2021, Ownership, learning, and beliefs, *Quarterly Journal of Economics* 136, 1665–1717.

Hirshleifer, David, Jun Li, and Jianfeng Yu, 2015, Asset pricing in production economies with extrapolative expectations, *Journal of Monetary Economics* 76, 87–106.

Jin, Lawrence J, and Pengfei Sui, 2022, Asset pricing with return extrapolation, *Journal of Financial Economics* 145, 273–295.

Johannes, Michael, Lars A Lochstoer, and Yiqun Mou, 2016, Learning about consumption dynamics, *Journal of Finance* 71, 551–600.

Khaw, Mel Win, Luminita Stevens, and Michael Woodford, 2017, Discrete adjustment to a changing environment: Experimental evidence, *Journal of Monetary Economics* 91, 88–103.

Koijen, Ralph SJ, and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Kuhnen, Camelia M, 2015, Asymmetric learning from financial information, *Journal of Finance* 70, 2029–2062.

Liu, Haoyang, and Christopher Palmer, 2021, Are stated expectations actual beliefs? New evidence for the beliefs channel of investment demand, NBER Working paper.

Moyer, Robert, and Thomas Landauer, 1967, Time required for judgments of numerical inequality, *Nature* 215, 1519–1520.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *Review of Financial Studies* 35, 2190–2245.

Nagel, Stefan, and Zhengyang Xu, 2023, Dynamics of subjective risk premia, *Journal of Financial Economics* forthcoming.

Veronesi, Pietro, 1999, Stock market overreactions to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.

Woodford, Michael, 2020, Modeling imprecision in perception, valuation, and choice, *Annual Review of Economics* 12, 579–601.

Internet Appendix for "Insensitive Investors"

CONSTANTIN CHARLES, CARY FRYDMAN, and METE KILIC[1]

[1]Citation format: Charles, Constantin, Cary Frydman, and Mete Kilic, Internet Appendix for "Insensitive Investors,", *Journal of Finance* [DOI String]. Please note: Wiley is not responsible for the content or functionality of any supporting information supplied by the authors. Any queries (other than missing material) should be directed to the authors of the article.

IA.1. Does the transmission strength vary over time?

In our conceptual framework, we assume that the prior variance, σ_p^2, and the signal variance, σ_ϵ^2, are constant in the time series. An immediate implication of these assumptions is that the transmission strength, x, is also constant over time.

It is plausible, however, that the cognitive signal precision may vary over time. For example, if subjects become fatigued over the course of the experiment, σ_ϵ^2, may be larger for elicitations that occur later in the experiment. Alternatively, by gaining experience with the task, cognitive noise may become smaller towards the end of the experiment, corresponding to a decrease in σ_ϵ^2. Thus, it is an empirical question as to whether x remains constant, increases, or decreases over time.

Here we investigate the stability of transmission strength by re-estimating our main regressions and allowing x to vary between the first and second half of our experiments. Table IA.1 shows that in all three experiments, we cannot reject the null hypothesis that x is constant over time. The dummy variable "Late" takes the value 1 if the elicitation is in the second half of the experiment, and it takes the value 0, otherwise. None of the interaction terms are significantly different from zero and they all have small magnitudes, suggesting that transmission strength does not vary significantly over time.

Table IA.1
Stability of Weak Transmission

Sample: p	Experiment 1 (1)	Experiment 2 (2)	Experiment 3 (3)
d	0.579***	0.838***	0.846***
	(0.075)	(0.071)	(0.063)
d x Late	0.019	0.059	0.025
	(0.085)	(0.078)	(0.065)
λ	−0.219**	−0.346***	−0.319***
	(0.101)	(0.111)	(0.122)
λ x Late	0.049	−0.167	−0.115
	(0.129)	(0.119)	(0.136)
Late	−0.116	−0.212	−0.075
	(0.417)	(0.381)	(0.313)
Constant	1.861***	0.662*	0.648**
	(0.369)	(0.349)	(0.306)
Observations	2,400	2,400	2,400

This table presents results from mixed effects regressions of log(WTP) (p) on log subjective expected payoff (d) and perceived volatility (λ), interacted with a dummy variable (Late) that is equal to one for elicitation periods 5-8. Columns 1, 2, and 3 use data from Experiments 1, 2, and 3, respectively. All regressions include a random effect for d and λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

IA.2. Screenshots of instructions and decision screens

A. Experiment 1

A.1. Instructions

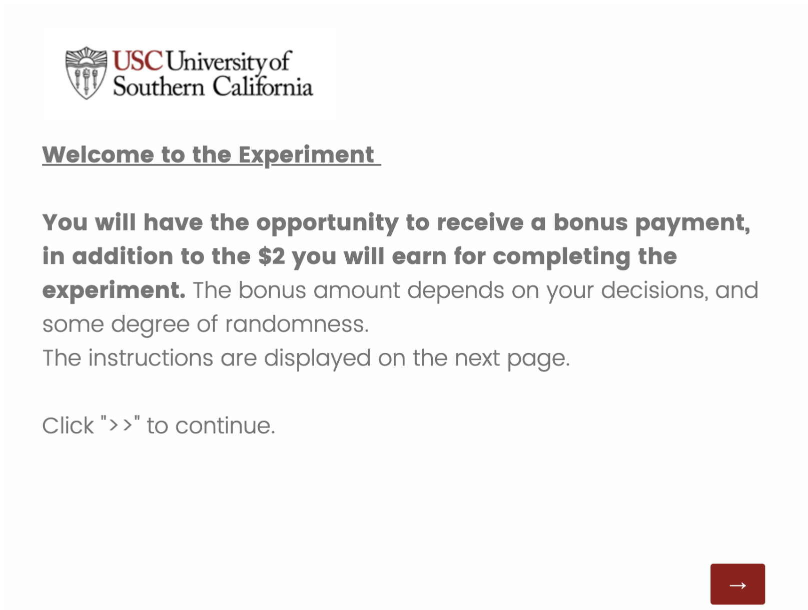
USC University of Southern California

Welcome to the Experiment

You will have the opportunity to receive a bonus payment, in addition to the $2 you will earn for completing the experiment. The bonus amount depends on your decisions, and some degree of randomness.
The instructions are displayed on the next page.

Click ">>" to continue.

→

Instructions

The experiment consists of 30 periods. In each period, there are two possible states: the black state and the yellow state. In the first period, we will select either the black state or the yellow state with 50% chance, but you will not see which state we selected. In every period, there is 20% chance that the state changes, and there is an 80% chance that the state remains the same as in the previous period. For example, in the first period, the state may be yellow, but it could switch to black a few periods later.

Whether the state is yellow or black is important, because it affects the dividend of the stock.

In each period, you will observe the **dividend of the stock**. The stock can pay out one of five possible dividends:

$60
$85
$115
$135
$150

The charts below show you the probability of each dividend, for each of the two states. While you won't know which state was selected for sure, you can learn which one is more likely by observing the dividends. For example, suppose you saw many $60 dividends. Because the yellow state has a higher chance of delivering a $60 dividend, then it is more likely that the yellow state was selected during those periods. Remember, however, that the state can change in each period.

USC University of Southern California

Instructions (continued)

You can earn a bonus payment of up to $3.

You will observe the stock's dividend for 30 periods. After some periods, we will ask you two questions about the stock. The first question asks you about your expectation of the stock's dividend in the next period. The second question asks you for your willingness to pay for an investment. After the experiment, we will randomly select one of your answers from this part to determine your bonus payment. If you are unsure about the correct answer, give us your best guess.

How your bonus is determined

The below paragraphs give details on exactly how the bonus is determined. We have designed the bonus payment so that reporting your true belief or true willingness to pay is optimal for earning the highest bonus. *If you do not want to read the details about how the bonus is computed, you can skip to the next page.*

After some periods, you will answer a question about how likely you think it is that each of the five possible dividends will occur in the next period. We will randomly select one of these questions, and from the selected question, one of your predictions will be randomly chosen to determine your bonus. You will receive $3 if your answer is correct. For every percentage point that your answer is further away from the correct value, we will subtract 3 cents from your bonus.

We will also ask you for your willingness to pay for a potential investment. If we randomly select one of these questions, your bonus is determined as follows. You will be endowed with $210 of experimental wealth, which can be used for
the investment. Depending on your willingness to pay for the investment, you might or might not be able to invest.

- If you are *not* able to invest: Your bonus will be one hundredth of your experimental wealth, that is $2.10.
- If you *are* able to invest: Your bonus will be higher if you make a profit, but lower if you make a loss. We will add profits and subtract losses from your experimental wealth. For example, if your profit is $30, we add this profit to your experimental wealth ($210 + $30 = $240), so will receive a bonus of $2.40. In contrast, if you make a loss of for example $20, we subtract this loss from your experimental wealth ($210 - $20 = $190), so you will receive a bonus of $1.90.

We will determine whether you are able to invest using a "Becker-DeGroot-Marschak mechanism". This mechanism is designed so that it is in your best interest to give us your true willingess to pay. It works as follows. After you give us your willingess to pay, we draw a random price between $60 and $150. If the price that we draw is equal to or smaller than your stated willingess to pay, you are forced to invest at the drawn price. If the number is larger than your stated willingness to pay, you do not invest.

Suppose your true willingess to pay is $100. There are three possible cases:

Case 1: You state your true willingess to pay of $100. If the price we draw is $100 or lower, you invest. Thus, you only invest if the price is equal to or below your true willingness to pay. You never pay more than your true willingness to pay.

Case 2: You state a willingess to pay that is lower than $100. Suppose you state a (false) willingness to pay of $70. You will invest if the price is $70 or lower. Thus, you miss out on all the opportunities where the price is between $70 and $100.

Case 3: You state a willingess to pay that is higher than $100. Suppose you state a (false) willingness to pay of $130. You will invest if the price is $130 or lower. Thus, when the price is between $100 and $130, you will be forced to invest even though the price is above your true willingness to pay.

<u>The bottom line:</u> It is in your best interest to always state your true willingness to pay.

[The distributions were displayed to subjects before the first dividend realization and in each elicitation period.

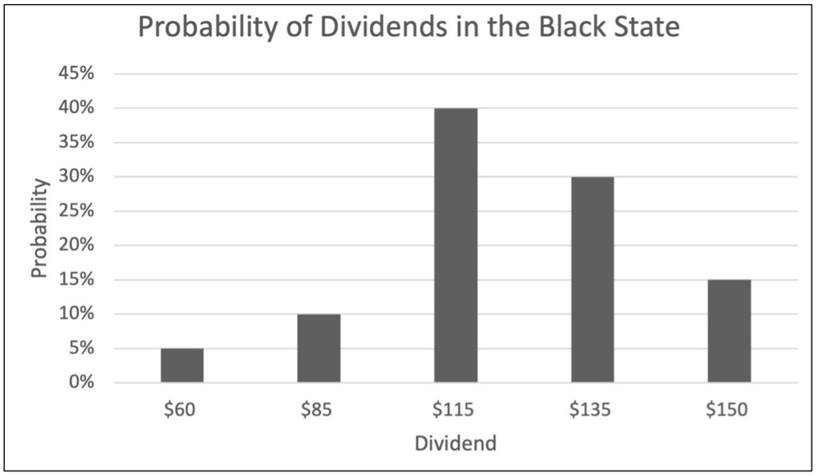

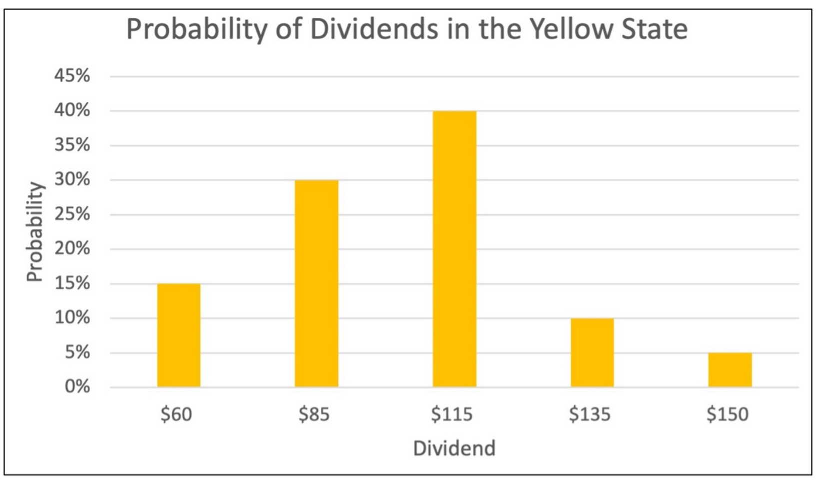

A.2. Decision screens

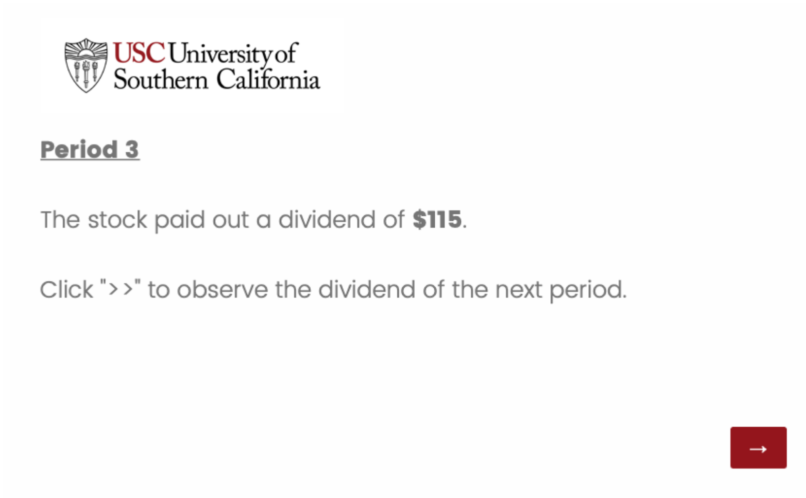

[Slider before initiation]

You have now observed the stock for 4 periods. The following summarizes the dividend in each period:

Period	Dividend
1	$135
2	$135
3	$115
4	$150

In this question, we would like to know your expectations of next period's dividend. Please let us know how likely you think it is that each dividend will occur in the next period.

Please type in the number to indicate the probability, in percent, that you attach to each scenario. The probabilities of the five scenarios have to sum up to 100%.

$150	0	%
$135	0	%
$115	0	%
$85	0	%
$60	0	%
Total	0	%

Suppose you could purchase the right to next period's dividend, before you knew how much it was worth. What is the highest price you'd be willing to pay now, for the right to receive next period's dividend?

Please use the slider to select your answer, in dollars.

60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150

Willingness to pay in $

[Slider after initiation]

You have now observed the stock for 4 periods. The following summarizes the dividend in each period:

Period	Dividend
1	$135
2	$135
3	$115
4	$150

In this question, we would like to know your expectations of next period's dividend. Please let us know how likely you think it is that each dividend will occur in the next period.

Please type in the number to indicate the probability, in percent, that you attach to each scenario. The probabilities of the five scenarios have to sum up to 100%.

$150	0	%
$135	0	%
$115	0	%
$85	0	%
$60	0	%
Total	0	%

Suppose you could purchase the right to next period's dividend, before you knew how much it was worth. What is the highest price you'd be willing to pay now, for the right to receive next period's dividend?

Please use the slider to select your answer, in dollars.

60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150

Willingness to pay in $

93.9

B. Experiment 2

B.1. Instructions

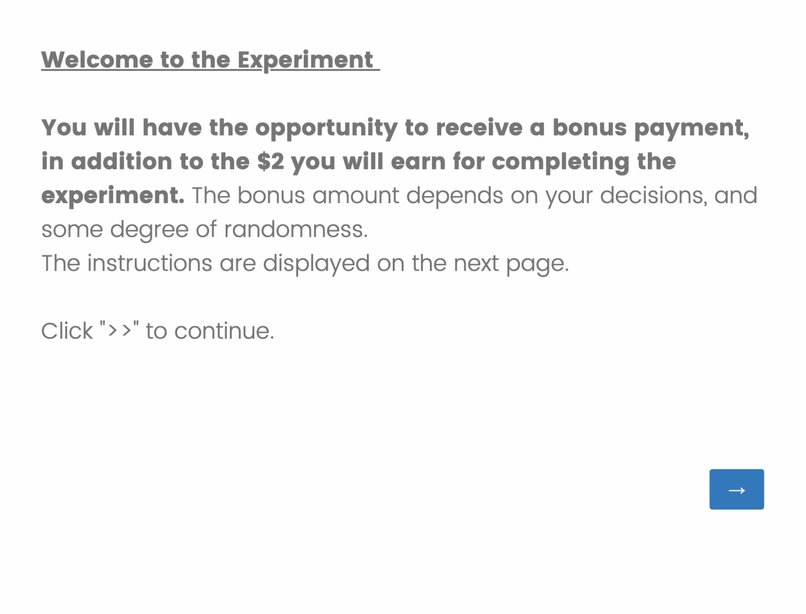
Welcome to the Experiment

You will have the opportunity to receive a bonus payment, in addition to the $2 you will earn for completing the experiment. The bonus amount depends on your decisions, and some degree of randomness.
The instructions are displayed on the next page.

Click ">>" to continue.

→

Instructions

The experiment consists of 8 periods. In each period, you will be presented with a new investment opportunity. Each investment opportunity can pay out one of five possible dividends:

$60
$85
$115
$135
$150

To help you evaluate each investment opportunity, you will receive a table that shows you the probability, in percent, of each dividend. Each investment opportunity is completely new and unrelated to any past or future investment opportunities. Your decisions about an investment opportunity do not affect which types of investment opportunities you will see in the future. You should evaluate each investment opportunity in isolation.

→

Instructions (continued)

You can earn a bonus payment of up to $1.50.

In each period, we will ask you for your willingness to pay for the investment opportunity. If you invest, the investment will pay out one of the five possible dividends according to the probabilities in the table. If a dividend has a higher probability, there is a higher chance that the investment will pay out that dividend. After the experiment, we will randomly select one of your answers to determine your bonus payment. If you are unsure about the correct answer, give us your best guess.

How your bonus is determined

The below paragraphs give details on exactly how the bonus is determined. We have designed the bonus payment so that reporting your true willingness to pay is optimal for earning the highest bonus. *If you do not want to read the details about how the bonus is computed, you can skip to the next page.*

In each period, we ask you for your willingness to pay for an investment opportunity. We will randomly select one of these questions to determine your bonus as follows. You will be endowed with $210 of experimental wealth, which can be used for the investment. With the investment, you can increase your experimental wealth if you make a profit, but you can also decrease your experimental wealth if you make a loss. At the end of the experiment, we will convert your experimental wealth into real dollars to calculate your bonus. For every two dollars of experimental wealth, you will receive one cent of bonus. For example, if your experimental wealth is $210, your bonus is $1.05.

Depending on your willingness to pay for the investment, you might or might not be able to invest.

- If you are *not* able to invest: Your experimental wealth will remain at $210.
- If you *are* able to invest: Your experimental wealth will be higher if you make a profit, but lower if you make a loss. We will add profits and subtract losses from your experimental wealth. For example, if your profit is $30, we add this profit to your experimental wealth ($210 + $30 = $240). In contrast, if you make a loss of for example $20, we subtract this loss from your experimental wealth ($210 - $20 = $190).

We will determine whether you are able to invest using a "Becker-DeGroot-Marschak mechanism". This mechanism is designed so that it is in your best interest to give us your true willingess to pay. It works as follows. After you give us your willingess to pay, we draw a random price between $60 and $150. If the price that we draw is equal to or smaller than your stated willingess to pay, you are forced to invest at the drawn price. If the number is larger than your stated willingness to pay, you do not invest.

Suppose your true willingess to pay is $100. There are three possible cases:

Case 1: You state your true willingess to pay of $100. If the price we draw is $100 or lower, you invest. Thus, you only invest if the price is equal to or below your true willingness to pay. You never pay more than your true willingness to pay.

Case 2: You state a willingess to pay that is lower than $100. Suppose you state a (false) willingness to pay of $70. You will invest if the price is $70 or lower. Thus, you miss out on all the opportunities where the price is between $70 and $100.

Case 3: You state a willingess to pay that is higher than $100. Suppose you state a (false) willingness to pay of $130. You will invest if the price is $130 or lower. Thus, when the price is between $100 and $130, you will be forced to invest even though the price is above your true willingness to pay.

The bottom line: **It is in your best interest to always state your true willingness to pay.**

B.2. Decision screens

[Slider before initiation]

Period 1

Dividend	Probability of Dividend
$60	10%
$85	20%
$115	30%
$135	25%
$150	15%
Total	**100%**

Suppose you could purchase the right to receive the dividend, before you knew how much it was worth. What is the highest price you'd be willing to pay now, for the right to receive the dividend?

Please use the slider to select your answer, in dollars.

60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150

Willingness to pay in $

[Slider after initiation]

Period 1

Dividend	Probability of Dividend
$60	10%
$85	20%
$115	30%
$135	25%
$150	15%
Total	**100%**

Suppose you could purchase the right to receive the dividend, before you knew how much it was worth. What is the highest price you'd be willing to pay now, for the right to receive the dividend?

Please use the slider to select your answer, in dollars.

60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150

Willingness to pay in $

104.7

C. Experiment 3

C.1. Instructions

Welcome to the Experiment

You will have the opportunity to receive a bonus payment, in addition to the $2 you will earn for completing the experiment. The bonus amount depends on your decisions, and some degree of randomness.
The instructions are displayed on the next page.

Click ">>" to continue.

→

Instructions

The experiment consists of 8 periods. In each period, you will be presented with a new investment opportunity. Each investment opportunity can pay out one of five possible dividends:

$60
$85
$115
$135
$150

Each investment opportunity is completely new and unrelated to any past or future investment opportunities. Your decisions about an investment opportunity do not affect which types of investment opportunities you will see in the future. You should evaluate each investment opportunity in isolation.

→

Instructions (continued)

You can earn a bonus payment of up to $3.00.

In each period, we will ask you twice for your willingness to pay for the investment opportunity. The first time, you will have no details about the investment opportunity. **We ask that you simply provide your best estimate for how much you'd be willing to pay for the investment opportunity, without seeing any information about it.**

The second time, you will receive a table that shows you the probability, in percent, of each dividend. The investment will pay out one of the five possible dividends according to the probabilities in the table. If a dividend has a higher probability, there is a higher chance that the investment will pay out that dividend. After the experiment, we will randomly select one of your answers to determine your bonus payment. If you are unsure about the correct answer, give us your best guess.

How your bonus is determined

The below paragraphs give details on exactly how the bonus is determined. We have designed the bonus payment so that reporting your true willingness to pay is optimal for earning the highest bonus. *If you do not want to read the details about how the bonus is computed, you can skip to the next page.*

In each period, we ask you for your willingness to pay for an investment opportunity. We will randomly select one of these questions to determine your bonus as follows. You will be endowed with $210 of experimental wealth, which can be used for the investment. With the investment, you can increase your experimental wealth if you make a profit, but you can also decrease your experimental wealth if you make a loss. At the end of the experiment, we will convert your experimental wealth into real dollars to calculate your bonus. For every dollar of experimental wealth, you will receive one cent of bonus. For example, if your experimental wealth is $210, your bonus is $2.10.

Depending on your willingness to pay for the investment, you might or might not be able to invest.

- If you are *not* able to invest: Your experimental wealth will remain at $210.
- If you *are* able to invest: Your experimental wealth will be higher if you make a profit, but lower if you make a loss. We will add profits and subtract losses from your experimental wealth. For example, if your profit is $30, we add this profit to your experimental wealth ($210 + $30 = $240). In contrast, if you make a loss of for example $20, we subtract this loss from your experimental wealth ($210 - $20 = $190).

We will determine whether you are able to invest using a "Becker-DeGroot-Marschak mechanism". This mechanism is designed so that it is in your best interest to give us your true willingess to pay. It works as follows. After you give us your willingess to pay, we draw a random price. If the price that we draw is equal to or smaller than your stated willingess to pay, you are forced to invest at the drawn price. If the number is larger than your stated willingness to pay, you do not invest.

Suppose your true willingess to pay is $100. There are three possible cases:

Case 1: You state your true willingess to pay of $100. If the price we draw is $100 or lower, you invest. Thus, you only invest if the price is equal to or below your true willingness to pay. You never pay more than your true willingness to pay.

Case 2: You state a willingess to pay that is lower than $100. Suppose you state a (false) willingness to pay of $70. You will invest if the price is $70 or lower. Thus, you miss out on all the opportunities where the price is between $70 and $100.

Case 3: You state a willingess to pay that is higher than $100. Suppose you state a (false) willingness to pay of $130. You will invest if the price is $130 or lower. Thus, when the price is between $100 and $130, you will be forced to invest even though the price is above your true willingness to pay.

The bottom line: **It is in your best interest to always state your true willingness to pay.**

C.2. Decision screens

[Slider before initiation]

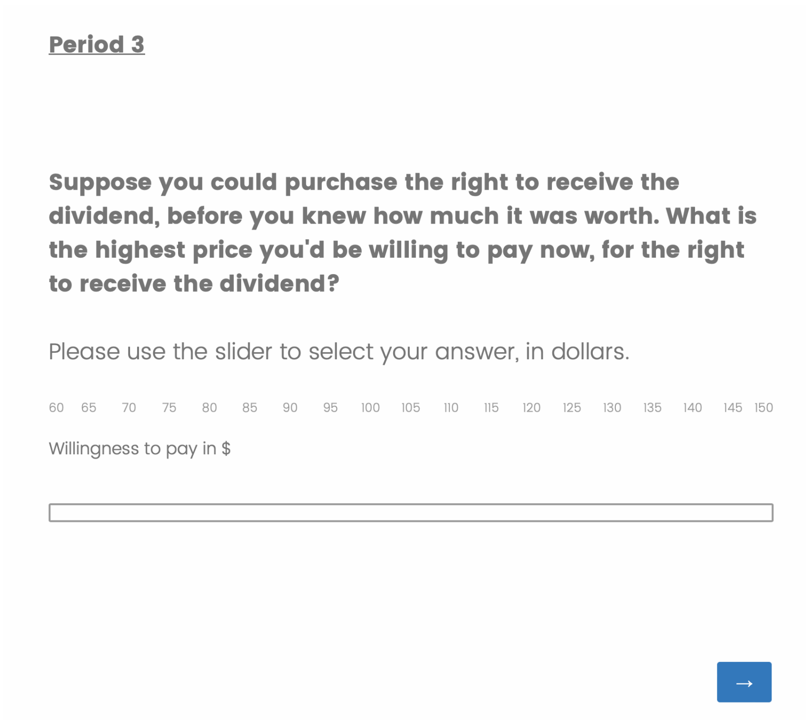

[Slider after initiation]

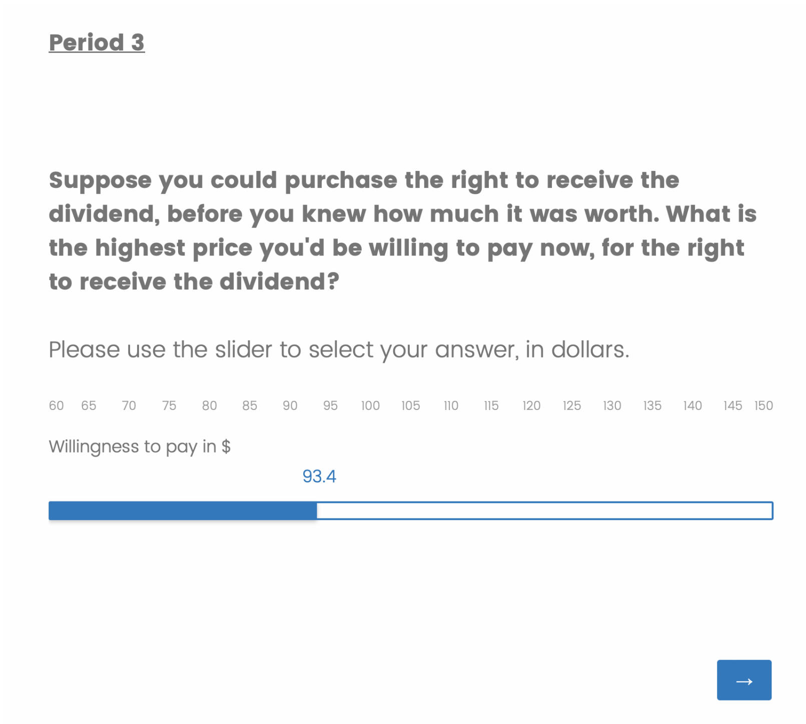

[Slider before initiation][2]

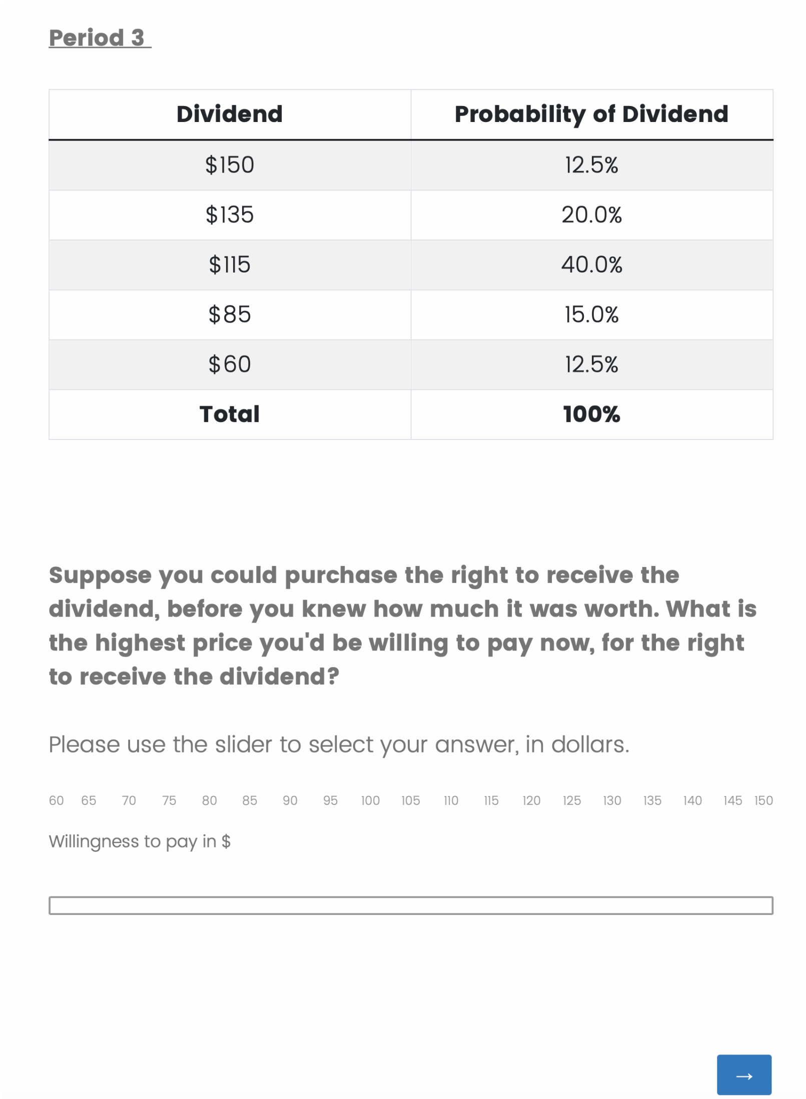

Period 3

Dividend	Probability of Dividend
$150	12.5%
$135	20.0%
$115	40.0%
$85	15.0%
$60	12.5%
Total	**100%**

Suppose you could purchase the right to receive the dividend, before you knew how much it was worth. What is the highest price you'd be willing to pay now, for the right to receive the dividend?

Please use the slider to select your answer, in dollars.

60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150

Willingness to pay in $

→

[2]The probabilities in Experiment 3 were displayed with one decimal, while the probabilities in Experiment 2 were displayed without decimals. We chose to include one decimal in Experiment 3 because some subjects in Experiment 1, from which the probability distributions are sourced, reported probabilities with one decimal. While we failed to display these decimals in Experiment 2, this occurred in only 25 of the 2,400 subject-rounds, which corresponds to about 1% of subject-rounds. There were no cases in which subjects reported probabilities with more than one decimal.

[Slider after initiation]

Period 3

Dividend	Probability of Dividend
$150	12.5%
$135	20.0%
$115	40.0%
$85	15.0%
$60	12.5%
Total	**100%**

Suppose you could purchase the right to receive the dividend, before you knew how much it was worth. What is the highest price you'd be willing to pay now, for the right to receive the dividend?

Please use the slider to select your answer, in dollars.

60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150

Willingness to pay in $

102.4

→

IA.3. Additional pre-registered analyses

In this section, we present additional results from our pre-registration of Experiment 1, available at https://aspredicted.org/6Z4_RLQ. First, in Figure IA.1, we show a binned scatter plot of $\mathbb{E}^b[D]/\mathbb{E}^*[D]$ and perceived disaster risk, using the subjective probability of the lowest payoff ($60) as the measure of perceived disaster risk. This is a measure of tail risk and is similar to the measure of disaster risk in Giglio et al. (2021a). Second, in Figure IA.2, we show a binned scatter plot of $\mathbb{E}^*[D]/P$ and perceived disaster risk. Finally, in Figure IA.3, we show a binned scatter plot of $\mathbb{E}^b[D]/P$ and perceived disaster risk.

Figure IA.1. Expected Cash Flows and Perceived Tail Risk

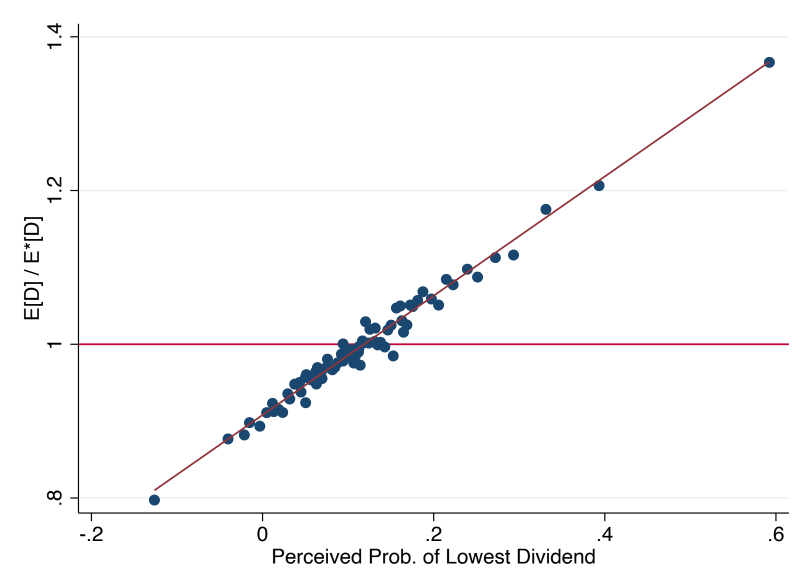

This figure is a binned scatter plot of $\mathbb{E}^b[D]/\mathbb{E}^*[D]$ and the perceived probability of the lowest payoff controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300. The data are from Experiment 1.

Figure IA.2. Subjective Expected Returns and Perceived Tail Risk

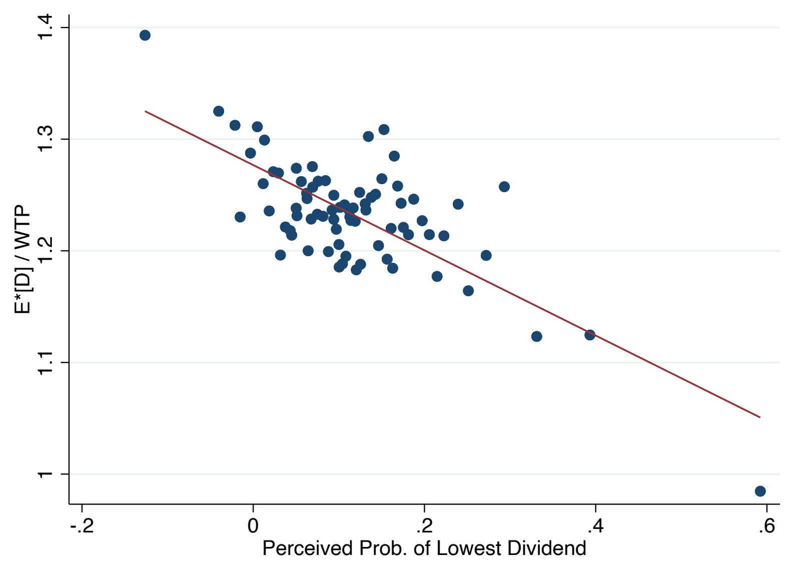

This figure is a binned scatter plot of $\mathbb{E}^*[D]/P$ and the perceived probability of the lowest payoff controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300. The data are from Experiment 1.

Figure IA.3. Objective Expected Returns and Perceived Tail Risk

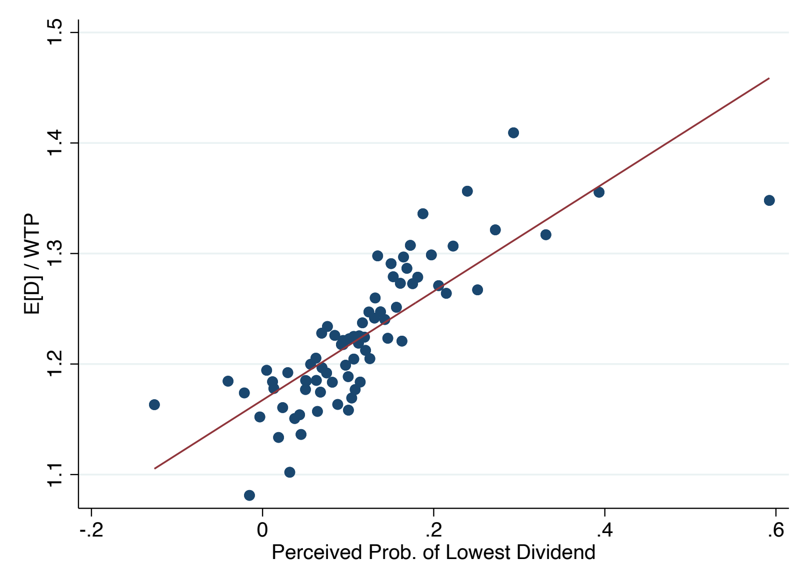

This figure is a binned scatter plot of $\mathbb{E}^b[D]/P$ and the perceived probability of the lowest payoff controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300. The data are from Experiment 1.

IA.4. Variation in the cognitive default (Experiment 3)

Here we present results from OLS fixed effects regressions, showing that the bulk of the variation in the cognitive default can be explained by a subject fixed effect.

Table IA.2

Variation in the Cognitive Default

$\bar{p}_t$	(1)	(2)	(3)
d_t			0.019
			(0.029)
d_{t-1}			0.005
			(0.034)
d_{t-2}			0.028
			(0.027)
p_{t-1}			0.088***
			(0.028)
p_{t-2}			0.038*
			(0.022)
Subject FE	Yes	Yes	Yes
Observations	2,400	1,800	1,800
R-squared	0.770	0.854	0.858

This table presents results from OLS fixed effects regressions of the log cognitive default ($\bar{p}$) on subject fixed effects. Column 1 uses all data from Experiment 3. Columns 2 and 3 restrict the sample to elicitation periods 3-8. Column 3 additionally includes current period and lagged log subjective expected payoff (d) and lagged log(WTP)(p) as explanatory variables. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

IA.5. Implications for the subjective risk-return relation

A. Omitted variable bias

A common finding that emerges from all three of our experiments is that there exists weak transmission from beliefs to valuations. Here, we highlight the implications of weak transmission for understanding properties of the subjective risk-return relation. We show theoretically, and provide supporting experimental evidence, that weak transmission will induce a strong bias in estimating the subjective risk-return relation. The basic problem that arises is as follows: even if an investor demands a higher risk premium as her perception of risk increases, failing to account for weak transmission can lead the econometrician to measure a negative subjective risk-return relation. This result, which we explain in more detail below, stems from the fact that weak transmission will induce a positive correlation between subjective expected payoffs and subjective expected returns.

Suppose the econometrician has data on the investor's WTP and the investor's subjective beliefs about D_{t+1}. In this scenario, it is straightforward to implement tests of the relation between perceived risk and subjective expected return shown in Equation (3) of the main text. That is, the econometrician can measure r_t as the difference between log expected payoff and log WTP as in (1). The econometrician can then regress the measured subjective r_t on λ_t, where the latter is also computed based on the investor's subjective beliefs about D_{t+1}. For any risk averse agent, there will be a positive relationship between perceived risk and subjective expected return, and the strength of this relationship is governed by the investor's risk aversion. However, this positive relationship need not hold if investors exhibit a weak transmission from beliefs to WTP.

If we plug WTP from (4) into the expected return in (1), we obtain

$$r_t = -(1-x)\bar{p}_t + (1-x)d_t + x\gamma\lambda_t - x\epsilon_t. \tag{IA.1}$$

When $x < 1$, the measured subjective expected return no longer depends only on the risk premium $\gamma\lambda_t$; instead, the measured subjective expected return will depend on *both* the risk premium and the expectation d_t.[3]

It follows that when there is weak transmission, the subjective expected return inferred from the agent's reported beliefs will differ from her risk-based discount rate. This difference becomes larger as the transmission becomes weaker (i.e., as $x \to 0$). Intuitively, when the reported payoff expectation d_t increases and p_t does not fully respond to the increase, the asset price becomes relatively "cheaper". This leads to a higher subjective expected return. Conversely, when an investor lowers her reported payoff expectation, she prices the asset lower, but not as low as under the frictionless benchmark. Thus, the weak transmission induces a positive correlation between subjective expected payoffs and subjective expected returns.

Equation (IA.1) implies that the weak transmission of beliefs to actions gives rise to an omitted variable bias in tests of the subjective risk-return relation. In particular, suppose that perceived risk λ_t and payoff expectation d_t are correlated and have an affine relation given by:

$$d_t = \alpha + \beta\lambda_t + \eta_t, \tag{IA.2}$$

where α and β are constants and η_t represents variation in d_t that is orthogonal to λ_t.

[3]When $x < 1$, the measured subjective expected return will also depend on the cognitive default $\bar{p}_t$. In Section C, we show that omitting the cognitive default from a regression of subjective expected return on perceived risk will work against the findings we present here.

Plugging (IA.2) into (IA.1), we obtain

$$\begin{aligned} r_t &= -(1-x)\bar{p_t} + (1-x)\alpha + [(1-x)\beta + x\gamma]\,\lambda_t + (1-x)\eta_t - x\epsilon_t, \\ &= -(1-x)\bar{p_t} + (1-x)\alpha + \left[\underbrace{(1-x)\beta - (1-x)\gamma}_{\text{bias}} + \gamma\right]\lambda_t + (1-x)\eta_t - x\epsilon_t \end{aligned} \tag{IA.3}$$

Hence, a univariate regression of r on λ generates a coefficient on risk equal to $(1-x)\beta - (1-x)\gamma + \gamma$. The term $(1-x)\beta - (1-x)\gamma$ represents the omitted variable bias, which can have a substantial effect on the estimated subjective risk-return relationship. The strength of the bias depends on the degree of the weak transmission (x), the loading of expected payoff on risk (β), and the price of risk (γ).

The coefficient on risk will be biased downward when the following conditions are met: (i) there is weak transmission ($0 < x < 1$) (ii) there is a negative correlation between expected payoff and risk ($\beta < 0$), and (iii) the price of risk is positive ($\gamma > 0$). For instance, suppose the asset moves between two states: a bad state with low expected payoffs and high risk, and a good state with high payoffs and low risk (as is the case in our experimental design). When the asset enters the bad state, the subjective expected return r_t increases because risk (λ_t) is higher (and γ is positive by assumption). But this effect is offset by the negative term $[(1-x)\beta - (1-x)\gamma]$, which creates the downward bias in estimation. Fixing β and γ, the downward bias becomes more severe as the transmission becomes weaker ($x \to 0$).

Figure IA.4 illustrates the role that weak transmission plays for the measured subjective risk-return relation in our experimental data. Each of the three panels plots the measured subjective risk-return relation for each of our three experiments. The x-axis in each panel is the conditional volatility of the dividend distribution, and the y-axis is the measured subjective expected return. Recall that we use the same exact belief distributions in each of the three experiments, and thus the x-axis in all panels is the same.

Figure IA.4. The Subjective Risk-Return Relation in the Three Experiments

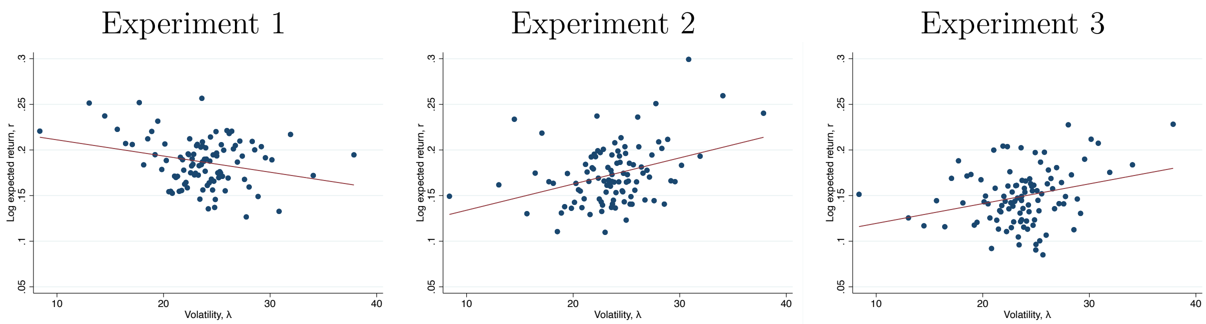

This figure shows binned scatter plots of subjective expected returns (r) and perceived risk (λ) in Experiments 1, 2, and 3 controlling for subject fixed effects. The sample size in each Panel is 2,400 and the number of subjects is 300.

In Panel A, we find a *negative* subjective risk-return relation. At first glance, this suggests that subjects discount the asset to a greater extent when there is *less* risk. However, our theoretical analysis above suggests an alternative interpretation: when $x < 1$ and expected payoffs and risk are negatively correlated (as they are in all experiments), the slope of the subjective risk-return relation is biased downward. If x is sufficiently low, this can lead to a negative subjective risk-return relation, even if subjects do apply greater discounts when there is more risk. The estimated x from Experiment 1 is relatively low at 0.61, and thus the weak transmission is likely responsible for the negative subjective risk-return relation.

Panels B and C provide the subjective risk-return relation for Experiments 2 and 3. These two panels paint a very different picture, as they reveal a positive relation between risk and measured subjective expected returns. Recall that the quantity of risk remains exactly the same as we move from Experiment 1 to Experiments 2 and 3. Indeed, the only variable in the bias term of Equation (IA.3) that changes across experiments is x. In particular, we find that x increases significantly from 0.61 in Experiment 1, to 0.82 in Experiment 2 and 0.83 in Experiment 3. This increase in x is sufficiently large that it flips the sign of the measured subjective risk-return relation.

Equation (IA.1) from our conceptual framework also provides a recipe for how we can restore the positive subjective risk-return relation in Experiment 1: by controlling for subjective expected payoffs (d) in a regression of subjective expected returns (r) on perceived risk (λ). Table IA.3 shows that the sign of the coefficient on risk flips from negative to positive when adding a control for subjects' reported expected payoffs in Experiment 1. By "adding back" the omitted variable (d) in Column 2, we can eliminate the bias in the coefficient on λ.[4]

Table IA.3

The Subjective Risk-Return Relation in Experiment 1

r	(1)	(2)
d		0.390^{***}
		(0.050)
λ	-0.173^{**}	0.195^{***}
	(0.078)	(0.073)
Constant	0.227^{***}	-1.699^{***}
	(0.020)	(0.242)
Observations	2,400	2,400

This table presents results from mixed effects regressions of subjective expected returns (r) on subjective expected payoff (d) and perceived volatility (λ). These regressions include a random effect for d and λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The data are from Experiment 1.

A further concern is that the relation between subjective payoff expectations and perceived risk, represented by β in Equation (IA.2), might vary over time. As shown above, this parameter plays a significant role in applications to the subjective risk-return relationship. In particular, as illustrated in Equation (IA.3), $\beta < 0$ is a sufficient condition for a downward bias in the risk-return relation if $\gamma > 0$ and $x < 1$.

[4] Suppose that $\bar{p}$ and λ are correlated through d. In this case, controlling for d will eliminate any bias in the coefficient on λ that results from omitting $\bar{p}$ (also see Section C below.)

Table IA.4 tests whether β varies over time by regressing d on λ, interacted with a dummy variable ("Late"). This dummy variable takes the value 1 if the elicitation is in the second half of the experiment, and it takes the value 0, otherwise. Note that the beliefs data are identical across the three experiments; so it is sufficient to check this result in the data generated in Experiment 1 only. We find that β remains negative in both halves of the experiment but it is closer to zero in the latter half, as shown in table below. However, the coefficient in the second half -0.875 + 0.547 = -0.328 is also statistically different from zero ($p < 0.001$).

Table IA.4
Variation in β

λ	-0.875^{***}
	(0.045)
λ x Late	0.547^{***}
	(0.066)
Late	-0.187^{***}
	(0.017)
Observations	2,400

This table presents results from mixed effects regressions of log expected payoff (d) on perceived volatility (λ), interacted with a dummy variable (Late) that is equal to one for elicitation periods 5-8. All regressions include a random effect for λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

B. Connection to Giglio et al. (2021a)

The survey conducted by Giglio et al. (2021a) is similar to our experimental paradigm, in that we also collect data on beliefs and actions at the individual level. As in Giglio et al. (2021a), we regress actions on expectations and find that the empirical link is weaker than predicted by frictionless models. It is also worth noting that, like us, Giglio et al. (2021a) find a negative relationship between subjective expected returns and perceived risk. Our explanation for this pattern (at least in Experiment 1), is driven by a combination of cognitive noise and omitted variable bias. We caution that such a mechanism cannot be used to justify the negative relationship between subjective expected returns and perceived risk that Giglio et al. (2021a) document. This is because our results rely on time series variation in beliefs and actions within an individual. The results in Giglio et al. (2021a) rely on cross-sectional variation in beliefs and actions, where all investors face the same equilibrium asset price and form heterogeneous subjective return expectations conditional on that price. Thus, while the insensitivity between actions and beliefs demonstrated in both studies may derive from a common mechanism of cognitive noise, our data cannot speak directly to the pattern of subjective expected returns uncovered by Giglio et al. (2021a).

C. The role of the cognitive default

The cognitive default $\bar{p}_t$ is not observable in Experiments 1 and 2. Here, we discuss how omitting $\bar{p}_t$ affects our estimates of the subjective risk-return relation in Experiments 1 and 2. As discussed in Section B.2, it is possible that $\bar{p}_t$ and d_t are positively correlated. This potential correlation along with the relation between d_t and λ_t in Equation (IA.2) may give rise to a correlation between $\bar{p}_t$ and λ_t. Such a correlation might lead to omitted variable bias in a univariate regression of r_t on λ_t.

In what follows, we show that this bias works against our findings of (i) a negative subjective risk-return relationship in Experiment 1 and (ii) a more positive subjective risk-return relationship in Experiments 2 and 3 compared to Experiment 1. In particular, consider the case where there is an affine relationship between $\bar{p}_t$ and d_t:

$$\bar{p}_t = a + bd_t, \tag{IA.4}$$

with $b > 0$. Plugging d_t from Equation (IA.2) into Equation (IA.4) implies

$$\bar{p}_t = a + b\alpha + b\beta\lambda_t + b\eta_t. \tag{IA.5}$$

The correlation between $\bar{p}_t$ and λ_t introduces a bias in the coefficient on λ_t in a univariate regression of r_t on λ_t. In particular, plugging Equation (IA.5) into Equation (IA.3) shows that this bias is equal to

$$-(1-x)b\beta. \tag{IA.6}$$

This bias is induced by the omission of $\bar{p}_t$ from the regression of r_t on λ_t. Note that β is negative and identical across all three experiments, and $0 < x < 1$ in all three experiments. Therefore, when $b > 0$, the coefficient on λ_t is biased upward in the regression of r_t on λ_t.

As a result, the omission of $\bar{p}_t$ works against our finding that the risk-return relation is negative in Experiment 1, and that it becomes positive in Experiments 2 and 3. The reason is that x is smaller in Experiment 1, which leads to a larger upward bias in Experiment 1. Further, since we elicit $\bar{p}_t$ in Experiment 3, we can check whether omitting $\bar{p}_t$ has an effect on the measured subjective risk-return relation (at least in Experiment 3). The following table shows that controlling for the cognitive default ($\bar{p}$) in a regression of subjective expected returns (r) on perceived risk (λ) does not materially change the coefficient on λ.

Table IA.5

The Cognitive Default and the Subjective Risk-Return Relation

r	(1)	(2)
$\bar{p}$		-0.232^{***}
		(0.031)
λ	0.208^{**}	0.240^{**}
	(0.096)	(0.094)
Constant	0.098^{***}	1.125^{***}
	(0.024)	(0.144)
Observations	2,400	2,400

This table presents results from mixed effects regressions of subjective expected returns (r) on the log cognitive default ($\bar{p}$) and perceived volatility (λ). These regressions include a random effect for $\bar{p}$ and λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses below the coefficient estimates. The coefficients and standard errors for λ are multiplied by 100. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The data are from Experiment 3.