Insensitive Investors*

Constantin Charles Cary Frydman Mete Kilic

June 29, 2022

Abstract

We show theoretically that the weak transmission of beliefs to actions induces a strong bias in basic asset pricing tests. In particular, expected returns can appear to decline in risk when investors weakly transmit their payoff expectations into willingness to pay. We experimentally test this prediction and find that subjects exhibit an extremely weak transmission of beliefs to actions, which generates a negative risk-return relation. We argue that the weak transmission is due to cognitive noise and demonstrate that cognitive noise causally affects the risk-return relation. Our results highlight the importance of incorporating weak transmission into belief-based asset pricing models.

*We are grateful to Nick Barberis, Alex Chinco, Ricardo De La O, David Hirshleifer, Lawrence Jin, Stijn van Nieuwerburgh, Selale Tuzel, Sunil Wahal, Mitch Warachka as well as participants at the Miami Behavioral Finance Conference, California Corporate Finance Conference, Arizona State University, Hong Kong University, City University of Hong Kong, and USC Marshall for helpful comments. All three authors are from Marshall School of Business, University of Southern California. Charles: ccharles@marshall.usc.edu. Frydman: cfrydman@marshall.usc.edu. Kilic: mkilic@marshall.usc.edu. Frydman acknowledges NSF grant #1749824 for financial support. This paper supersedes our previous work titled "Discounting Less in Bad Times: Shining the Light on Cash Flow Expectations".

1 Introduction

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 (2022), Nagel and Xu (2022b)). 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 et al. (2015), Barberis et al. (2018), Bordalo et al. (2019), Jin and Sui (2022), Nagel and Xu (2022a)). 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 beliefs 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 et al. (2017), Ameriks et al. (2020), Liu and Palmer (2021), Beutel and Weber (2022)). 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)).

In this paper, we analyze and experimentally test how the weak transmission of subjective beliefs to actions affects the basic building blocks of asset pricing. In the theoretical part of the paper, we show that failing to account for a weak transmission of beliefs to actions can fundamentally alter the interpretation of the risk-return relationship. Our main theoretical result is that inference from regressions of subjective expected returns on perceived risk is severely biased when investors do not fully transmit their beliefs into actions. Intuitively, when an investor raises her subjective expected payoff, the weak transmission dampens her associated increase in willingness to pay. The larger increase in expected payoff compared to willingness to pay leads to an increase in the subjective expected return. The weak

transmission therefore induces a positive correlation between expected returns and expected payoffs. Thus, if the econometrician runs a univariate regression of expected return on perceived risk, there will be an omitted variable bias when expected payoff correlates with perceived risk. If the link between beliefs and actions is weak enough, the measured risk-return relation can become negative. Importantly, controlling for the omitted expected payoff variable will restore the positive risk-return relation.

We test our theoretical predictions across two controlled experiments. We design our first experiment (Experiment 1) to provide three main advantages that complement data from surveys. First, we exogenously set the payoff process and control the subject's information set; we can then make quantitative statements about how subjective expected returns differ from objective (statistical) expected returns. Second, we incentivize subjects to price a one-period dividend strip in a partial equilibrium setting. For each subject, we elicit both their full distribution of beliefs about next period's stochastic dividend and their willingness to pay (WTP) for the dividend. The partial equilibrium aspect allows us to study the relationship between expectations and valuations at the subject level, without requiring the subject to be the marginal investor. Third, we back out the implied subjective expected return from valuations and expectations. This method circumvents common concerns in the survey literature about respondents not understanding what is meant by "expected returns" (Cochrane (2011), Cochrane (2017)).

We first test for the weak transmission of subjective payoff expectations to WTP, which is a necessary assumption for our theoretical predictions. In a frictionless model, a one unit increase in a subject's expected payoff generates a one unit increase in her WTP. Our experimental data strongly depart from this frictionless benchmark: we find that a one unit increase in expected payoff leads to only a 63% increase in WTP. Importantly, our experimental design shuts down all institutional frictions that can plausibly explain the disconnect between beliefs and actions in the field – such as costly portfolio monitoring, capital gains taxes, default retirement contributions, and leverage and short-selling constraints. 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.

After establishing the weak pass-through of beliefs to WTP in our experimental setting, we turn to testing our main predictions about the risk-return relation. When estimating a simple regression of measured expected returns on perceived risk, we find a strong *negative* relationship. This result is striking given that the average subject in our experiment is risk averse. Our conceptual framework provides an explanation for the negative relationship between expected returns and perceived risk. By omitting the expected payoff from the regression, there is a severe downward bias in the risk-return relation. Importantly, when we add expected payoff to the regression, the risk-return relation flips sign to become positive. This result suggests that it is important to account for the weak transmission, especially in simple regressions of returns on risk.

Given the absence of any institutional frictions in our experiment, we argue that a psychological explanation is responsible for the observed weak transmission of beliefs to WTP. We interpret the weak transmission through the lens of a new agenda in behavioral economics which argues that the decision-making process is subject to inherent *cognitive noise* (see Woodford (2020) for a review). The noise arises in the investor's mind due to cognitive constraints, and it increases with the complexity of the task at hand. Crucially, the noise leads to systematic decision biases: the investor is aware of this noise, and consequently shades her decision toward a "default value" that does not vary with the specific problem at hand (Enke and Graeber (2021)).

For example, when coming up with the valuation for an asset, an investor will naturally lean on her beliefs about future payoffs. But she may be uncertain about how to arrive at an appropriate valuation given these beliefs and her risk appetite. As such, she selects a valuation that is somewhere between the one dictated by her stated beliefs and a default valuation. The compression of actions towards a 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 shading of valuations towards a cognitive default immediately dampens the transmission of stated beliefs to WTP.

In our second experiment (Experiment 2), we manipulate the level of cognitive noise to assess its impact on the degree of weak transmission and the risk-return relationship. To do

so, we draw on the finding from Enke and Graeber (2021) that subjects report higher levels of cognitive uncertainty 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 dividends 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 dividends, 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. By comparing the sensitivity of WTP to beliefs across experiments, we can assess 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 87%, compared to 63% in Experiment 1. 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.[1]

We then test how the greater pass-through of beliefs to actions affects the risk-return relation. It is worth emphasizing that any change in the risk-return relation that we find across experiments must be due to the increased transmission since we hold constant all other

[1]Our finding is similar to, but distinct from, the experimental result in Hartzmark et al. (2021) where subjects react more strongly to information about goods that they own compared to those that they do not own. In Hartzmark et al. (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.

parameters. As predicted by our theory, we find that the slope of the measured risk-return relation increases significantly compared to Experiment 1. The increased pass-through from beliefs to actions is so much stronger in Experiment 2 that it flips the sign and restores a positive risk-return relation – even without controlling for expected payoff.

Our results suggest that cognitive noise is a key source of the disconnect between beliefs and actions. Because the noise arises inside the investor's mind, it is distinct from classical measurement error in surveys. If measurement error was solely responsible for the weak transmission, we would not expect any change in transmission strength across our two experiments, yet we find a substantial difference. The distinction between the two mechanisms is important because cognitive noise and measurement error will have different effects on aggregation, as behavior with cognitive noise depends heavily on a default value which can differ across investors (Liu and Palmer (2021)).

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, and it can generate a wedge between subjective and objective expected returns. Relatedly, Nagel and Xu (2022b) document a systematic difference between the cyclical behavior of subjective and objective expected returns; our framework suggests that weak transmission may be one potential explanation for this pattern in the data. While survey data is clearly valuable for unveiling differences in subjective and objective expectations, our work suggests caution in quantitative modeling approaches that assume agents fully act on their beliefs. At the same time, adding a weak transmission channel to existing models may provide an opportunity to further improve quantitative fits.

We also contribute to a burgeoning literature studying the weak transmission of beliefs to actions in the field. In a related paper, Liu and Palmer (2021) find that the beliefs investors report on surveys do not contain all the information used in actual investment decisions. Respondents who are less confident in their beliefs tend to rely on factors – such as past returns – to a greater extent in their decision making. Our results are wholly consistent with

these effects, but we additionally provide causal evidence in a controlled setting. Beutel and Weber (2022) study the causal effects of survey beliefs on portfolio choices. They also find a low sensitivity of portfolios to beliefs, but argue the primary cause is due to leverage and short sale constraints. In our experiments, institutional constraints are shut down by design, which suggests a substantial portion of the disconnect is driven by cognitive factors. Along these lines, Barberis and Jin (2022) 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. Their explanation relies on a "model free" system of decision-making, which is disconnected from the "model based" system that generates the beliefs reported by investors. 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 external factors such as measurement error or institutional constraints.

The rest of this paper is organized as follows. Section 2 presents a conceptual framework that illustrates the impact of the weak transmission of beliefs to WTP on the risk-return relation. Sections 3 and 4 present the results from Experiment 1 and Experiment 2, respectively. Section 5 discusses how our results relate to field evidence and broader implications for asset pricing. Section 6 concludes with directions for future work.

2 Conceptual framework

We start 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 pricing and the measurement of the risk-return relation.

2.1 Frictionless benchmark

Suppose that an agent can invest in an asset which delivers a stochastic payoff D_t at each time t. 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}.[2] 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, expected returns, and expected payoffs in logs which will simplify the predictions that we derive here and test in the next section.

The expected return r_t is equivalent to the discount rate that the agent applies to the payoff d_t, in order to generate her WTP, p_t. This equivalence is easily seen by rearranging (1) into $p_t = d_t - r_t$. We assume that the agent discounts the expected payoff based on her perceived riskiness of the payoff and her risk preference represented by

$$r_t = \gamma \lambda_t, \tag{2}$$

where γ is the price of risk (e.g., risk aversion) and λ_t is the quantity of risk implied by the agent's subjective beliefs about d_t (e.g., conditional volatility).[3] Hence, the expected return r_t represents compensation for risk which is a common notion in asset pricing for undiversifiable risks.[4]

If the econometrician has data on the investor's subjective distribution of D_{t+1} and her WTP, p_t, then it is straightforward to implement tests of the relation between risk and expected return shown in (2). That is, the econometrician can measure r_t as the difference between expected payoff and WTP as in (1). The econometrician can then regress the measured r_t on λ_t, where the latter is also computed based on the investor's subjective distribution of D_{t+1}. For any risk averse agent, there is a positive relationship between risk

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

[3] 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 necessity for time discounting, which will be the case in our experimental design.

[4] For instance, Cochrane (2011) considers "discount rate" and "expected return" to be equivalent in his discussion of time variation in discount rates.

and expected return, and the strength of this relationship is governed by the investor's risk aversion.

2.2 Insensitive actions

2.2.1 Compression towards a default valuation

We have so far maintained the standard assumption that the investor values the asset by assessing the expected payoff (from her subjective beliefs) and then applying a discount based on perceived risk. In particular, her beliefs "pass through" to her WTP in a 1-1 fashion, such that a one unit increase in her subjective expected payoff generates a one unit increase in the price she is willing to pay, controlling for perceived risk. Here we relax this assumption and 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)).

In particular, 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: it is the price the agent arrives at when she has no uncertainty about her beliefs, her risk aversion, or how to optimally combine these components. 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 (2021)). 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 what her true valuation is.

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 still arises in the decision process that transforms precise expectations to actions.

Following Gabaix (2019) and Enke and Graeber (2021) 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 $N(\bar{p}, \sigma_p^2)$. Here $\bar{p}$ is a "cognitive default" which represents the average valuation in a similar class of problems. It is the valuation she would choose before drawing her noisy cognitive signal. The agent then combines her prior and signal to come up with the posterior mean, which is the WTP that she reports:

$$\begin{aligned} p_t &= (1-x)\bar{p} + xp_t^0 \\ &= (1-x)\bar{p} + xd_t - x\gamma\lambda_t + x\epsilon_t \end{aligned} \tag{3}$$

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. See Appendix B.1 for a derivation. Importantly, a one unit increase in d_t now leads to an increase in p_t by only x units.

To help illuminate the mapping between our framework and applications, consider an investor who is assessing her valuation for 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, which could be yesterday's price or some weighted average price from the recent past. The default is very likely to depend on an investor's past experience with the asset, and thus can differ across investors (Liu and Palmer (2021)). The default price is represented by $\bar{p}$ in our framework. If enough investors behave in this manner, then the price of the stock market will adjust in the right direction, but not by enough.

It is important to point out that Equation (3) implies that WTP will also 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.

We emphasize that while cognitive noise readily generates a stickiness in actions, in gen-

eral, there are other factors that can also lead to an insensitivity between reported beliefs and actions. For example, in the field, default contribution rates or capital gains taxes may lead investors to act less aggressively than is dictated by their beliefs alone. Alternative behavioral models based on inattention or memory constraints may also drive a wedge between reported beliefs and actions. In what follows, the key assumption we rely on is encoded in Equation (3): agents underreact to their reported beliefs. The microfoundation based on cognitive noise helps us structure the predictions for our experiment, but the implications we now present apply more generally in the presence of alternative frictions.

2.2.2 Implications for the risk-return relation

The weak transmission of beliefs to WTP has important implications for the risk-return relation. The economic interpretation of the *measured* expected return $r_t = d_t - p_t$ changes in the case of weak transmission of beliefs to WTP. Plugging WTP from (3) into the expected return in (1), we obtain

$$r_t = -(1-x)\bar{p} + (1-x)d_t + x\gamma\lambda_t - x\epsilon_t. \tag{4}$$

The measured expected return with $x < 1$ no longer depends only on the risk premium $\gamma\lambda_t$ but also on the belief d_t.

In the case of weak transmission, the subjective expected return based on the agent's reported beliefs differs 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 expected return. Conversely, when an investor lowers her reported expectation, she prices the asset lower, but not as low as under the frictionless benchmark. Thus, the weak transmission induces a positive correlation between expected payoffs and expected returns. In the frictionless case, equation (4) becomes $r = \gamma\lambda_t$, and the expected return again corresponds to the risk premium only.

Equation (4) implies that the weak transmission of beliefs to actions gives rise to an omitted variable bias in tests of the 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:

$$d_t = \alpha + \beta\lambda_t + \eta_t, \tag{5}$$

where α and β are constants and η_t represents variation in d_t that is orthogonal to λ_t. Plugging (5) into (4), we obtain

$$\begin{aligned} r_t &= -(1-x)\bar{p} + (1-x)\alpha + \left[(1-x)\beta + x\gamma\right]\lambda_t + (1-x)\eta_t - x\epsilon_t, \\ &= -(1-x)\bar{p} + (1-x)\alpha + \left[\underbrace{(1-x)\beta - (1-x)\gamma}_{\text{bias}} + \gamma\right]\lambda_t + (1-x)\eta_t - x\epsilon_t \end{aligned} \tag{6}$$

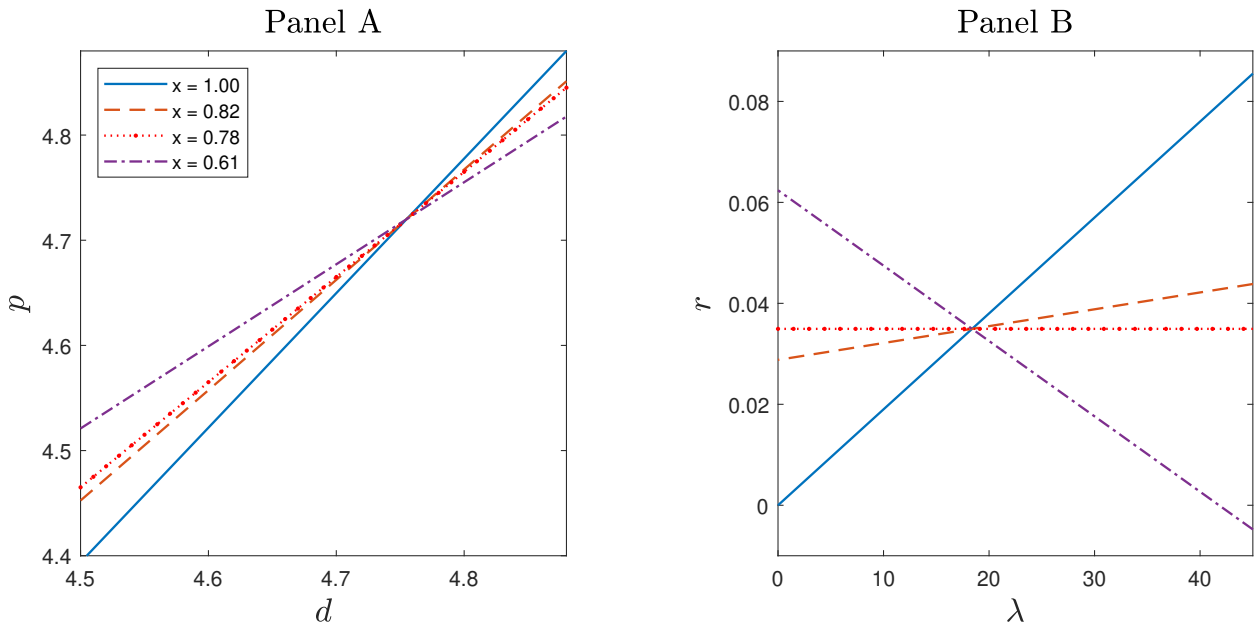
Hence, a univariate regression of r on λ generates a coefficient on risk equal to $(1-x)\beta - (1-x)\gamma + \gamma$. The term $(1-x)\beta - (1-x)\gamma$ represents the omitted variable bias, which can have a substantial effect on the estimated risk-return relationship. The strength of the bias depends 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 will be the case in our experimental design). When the asset enters the bad state, the 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 1 illustrates the consequences of the weak transmission on the risk-return relation. To create this figure, we fix the price of risk γ, the default value $\bar{p}$, and the parameters which govern the relation between payoff and risk (α, β). The only parameter we vary is x, which captures the strength of the transmission of beliefs to WTP.[5] When $x = 1$ (blue solid line),

[5]To create Figure 1, we pick parameter values based on our experimental results. In particular, γ is based on the coefficient on λ_t in a regression of p_t on d_t and λ_t, and (α, β) are based on a regression of d_t on λ_t. We set the value of $\bar{p}$ to the mean log payoff expectation.

Figure 1. The Impact of the Weak Transmission on the Risk-Return Relation



Notes: The figure illustrates numerical examples for different values of x. Panel A plots p against d using equation (3). Panel B plots r against λ using equation (6) where λ-d pairs are based on equation (5) with $\eta = 0$. The parameter values are $\bar{p} = 4.72, \gamma = 0.0019, \beta = -0.0068, \alpha = 4.88$.

the reported belief, d, is transmitted 1-1 into the WTP: p is therefore equal to d discounted by the risk premium, $\gamma\lambda$. However, as x decreases, p become less responsive to d (Panel A). Crucially, the slope of the risk-return relation also decreases as x decreases, despite holding constant the price and quantity of risk (Panel B). Equation (6) implies that there is a threshold value for x (equal to $\frac{\beta}{\beta-\gamma}$) for which risk and return will exhibit zero correlation (red dotted line). For any value of x less than this threshold, the risk-return relation flips sign and becomes negative (e.g., the purple dash-dotted line).

2.3 Predictions

Here we summarize the main implications of our conceptual framework and develop two testable predictions that we take to the experimental data. Our first prediction is on the measured risk-return relation in the presence of weak transmission.

Prediction 1. *If (i) the transmission of beliefs to willingness to pay is weak and (ii) payoff*

expectations and risk are negatively correlated, then regressing expected return on risk will lead to a downward biased coefficient on risk. The coefficient on risk increases if the payoff expectation is included as a control in the regression.

Because the amount of bias in the risk-return relation depends on the parameter x, it follows from our conceptual framework that manipulating x will systematically affect the risk-return relation. Because one source of weak transmission can be cognitive noise, we predict that manipulating cognitive noise will systematically affect the risk-return relation.

Prediction 2. *If payoff expectations and risk are negatively correlated, then exogenously decreasing cognitive noise will (i) strengthen the sensitivity of WTP to payoff expectations, (ii) strengthen the sensitivity of WTP to perceived risk, and (iii) increase the coefficient on risk in a regression of expected returns on risk.*

In the next section, we test these predictions in an experimental setting.

3 Experiment 1

3.1 Experimental design

3.1.1 Experimental setup

The goal of our experiment is to cleanly test for the weak transmission of beliefs to WTP and its implications for the risk-return relation. Importantly, our design shuts down several factors that Giglio et al. (2021a) suggest can generate a low sensitivity of portfolios 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 expected returns. We also infer the expected return from a subject's (i) WTP for the asset and (ii) reported beliefs about the asset payoff.

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.[6] 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:

$$\Pr(D_t|s_t = bad) \equiv (\$60, 0.15; \$85, 0.30; \$115, 0.40; \$135, 0.10; \$150, 0.05) \tag{7}$$

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

$$\Pr(D_t|s_t = good) \equiv (\$60, 0.05; \$85, 0.10; \$115, 0.40; \$135, 0.30; \$150, 0.15). \tag{8}$$

The distribution in the *good* state has a higher mean and lower volatility, compared with the dividend distribution in the *bad* state. We initialize the state in period 1 to be either *good* or *bad* with equal probability: $\Pr(s_1 = good) = 50\%$. 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 for two main reasons. First, the Markov chain induces substantial time series variation in the expected dividend. Moreover, the two-state switching process 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). Second, the switching process induces a negative time series correlation between the conditional expectation and volatility of the payoff. This negative correlation is important because it is one of the necessary conditions for Prediction 1 to obtain.[7] To ease comparability of behavior across subjects, we use the same realized sequence of thirty

[6] 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.

[7] The predictions in our conceptual framework rest on the assumption that the *subjective* expected payoff d_t and the perceived risk λ_t are negatively correlated (i.e., that $\beta < 0$). By definition, it is impossible to impose this correlation on subjective beliefs. However, we can increase the chance of observing such a correlation by imposing a negative correlation for a Bayesian agent. In Section 3.2 we confirm that the subjective expected payoff and perceived risk are indeed negatively correlated.

dividends for all subjects.

In 8 randomly chosen periods, we elicit a subject's full distribution of expectations about next period's dividend. In the remaining 22 periods, we do not elicit any expectations, and subjects simply observe the realized dividend.[8] 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 across subjects, and we ensure 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 distribution.

Importantly, we incentivize the expectations question and the WTP question. When we elicit a subject's distribution of 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 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 *good* and *bad* states depicted in equations (7) and (8). 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:

$$\begin{aligned} & q_t(q_{t-1}, D_t) \\ &= \frac{\Pr(D_t|s_t = bad)\Pr(s_t = bad|q_{t-1})}{\Pr(D_t|s_t = bad)\Pr(s_t = bad|q_{t-1}) + \Pr(D_t|s_t = good)\Pr(s_t = good|q_{t-1})} \\ &= \frac{\Pr(D_t|s_t = bad)(0.8q_{t-1} + 0.2(1 - q_{t-1}))}{\Pr(D_t|s_t = bad)(0.8q_{t-1} + 0.2(1 - q_{t-1})) + \Pr(D_t|s_t = good)(0.2q_{t-1} + 0.8(1 - q_{t-1}))}, \end{aligned} \tag{9}$$

where the expressions $\Pr(D_t|s_t = bad)$ and $\Pr(D_t|s_t = good)$ are defined in equations (7) and (8) (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: $\mathbb{E}[D_t|q_t] = q_t\mathbb{E}[D_t|s_t = bad] + (1 - q_t)\mathbb{E}[D_t|s_t = good]$. Similarly, the probability of

[8] We elicit beliefs in the same 8 (randomly chosen) periods for all subjects. See Internet Appendix IA.1 for screenshots of the experiment.

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, $\Pr(D_t = \$60|q_t) = q_t \Pr(D_t = \$60|s_t = bad) + (1 - q_t) \Pr(D_t = \$60|s_t = good)$.

The calculations above establish the Bayesian benchmark, which we use to incentivize subjects when they report their beliefs. In particular, we randomly pick one of the eight periods in which we elicit the distribution of beliefs and the WTP, and then pay subjects based on either the distribution question or the WTP question. If the distribution question is randomly chosen, then we randomly select one of the outcomes of the distribution 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 report their true 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 (9), our tests do not rely on subjects' ability to accurately compute q_t. Importantly, our conceptual framework in Section 2 does not depend on whether beliefs depart from Bayesian rationality. Indeed, the implications from Section 2 are based on reported beliefs, which are inherently subjective. 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.

3.1.2 Experimental procedures

We recruit 300 subjects from the online data collection platform, Prolific. The sample size and exclusion criteria are pre-registered on Aspredicted.org.[9] Subjects received $2 for 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.

3.2 Experimental results

3.2.1 Summary statistics

Our experiment with 300 subjects produces a panel dataset with 2,400 total observations (8 elicitations per subject). Table 1 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 payoffs (dividends), the time series of the Bayesian distribution is identical across subjects.

Table 1 reveals that subjects are quite accurate about the expected payoff on average, but there are sizeable deviations 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. Thus, on average, subjective expected returns are greater than 1. Because we elicit subjects' beliefs about the entire payoff distribution, we can measure perceived risk as the volatility of the elicited distribution. While the average and median perceived volatility are close to the Bayesian benchmark, perceived risk exhibits a lot more variation.

For all subsequent empirical analyses, we transform WTP and expectations to logs in

[9] 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 Appendix D for additional pre-registered analyses. Our analyses in Experiment 1 provide crucial motivation for the design of Experiment 2, which we also pre-register; details are provided in Section 4.

Table 1
Summary Statistics

		Mean	p25	p50	p75	SD	Min	Max
Subjective expected payoff	$\mathbb{E}^*[D]$	112.60	105.5	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
Subjective expected return	$\mathbb{E}^*[R] = \mathbb{E}^*[D]/P$	1.23	1.04	1.19	1.38	0.27	0.52	2.23
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

Notes: This table presents summary statistics for the main variables in our sample. The sample consists of 300 subjects, each elicited at 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.

order to be consistent with the conceptual framework in Section 2. In particular, p denotes $\log P$, d denotes $\log \mathbb{E}^*[D]$, λ denotes $\text{Vol}^*[D]$, and r denotes $\log \mathbb{E}^*[R]$. In addition, most of our analyses in the following sections focus on the within-subject time series relation between variables. Appendix E shows that there is also significant time-invariant variation in beliefs and WTP across subjects, a finding that is consistent with the field evidence in Giglio et al. (2021a).

3.2.2 Pricing results

We begin by testing the underlying assumption of our theory from Section 2, namely, the existence of weak transmission of beliefs to actions. Because our theory is about pricing and beliefs at the individual level, we allow for heterogeneity in parameters across subjects when conducting our empirical tests. We assume that the default WTP ($\bar{p}$), price of risk (γ), degree of weak transmission (x), and the loading of payoff expectations on perceived risk (β) are fixed 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 2 reports that the sensitivity of WTP to payoff expectations is 0.634, which is significantly below one. Figure 2 illustrates this result by showing that the best fitting line is much shallower than the 45-degree line. Of course, our univariate regression of

Table 2
Willingness to Pay, Subjective Expected Payoffs, and Perceived Risk

p	(1)	(2)
d	0.634***	0.610***
	(0.049)	(0.050)
λ		−0.195***
		(0.073)
Observations	2,400	2,400

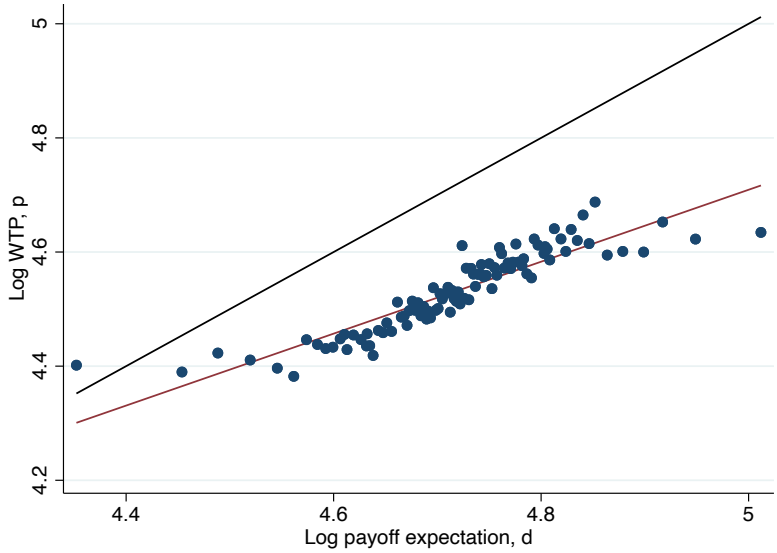
Notes: This table presents results from mixed effects regressions of WTP (p) on subjective expected payoffs (d) and perceived volatility (λ). These regressions include a random effect for d and λ, as well as for the intercept. Standard errors are clustered at the subject level and displayed in parentheses 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.

p_t on d_t has an omitted variable, namely, risk. In Appendix B.2, we show that the omission of λ_t 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 payoff expectations d_t and perceived risk λ_t is negative.[10] It follows that our underreaction result is robust to alternative measures of perceived risk besides volatility, since omitting the appropriate measure of risk will lead to an upward bias in the sensitivity of WTP to expected payoff. We show in Column 2 that the responsiveness of WTP to beliefs remains significantly below one after controlling for our assumed measure of risk, namely, conditional volatility. 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 payoff expectation constant.

Having established that $x < 1$, which is a necessary assumption for our predictions, we now turn to testing Prediction 1. Recall that this prediction states that if $x < 1$ and if payoff expectations and risk are negatively correlated, then regressing expected return on risk will lead to a downward biased coefficient on risk. Specifically, Equation (6) shows that the slope coefficient from regressing expected return on risk is given by $(1-x)\beta-(1-x)\gamma+\gamma$, where $(1-x)\beta-(1-x)\gamma$ is the omitted variable bias. We have already shown that $x < 1$ and

[10]By design, the correlation between payoff expectations and volatility is negative for a Bayesian investor. Figure A.1 demonstrates that this negative correlation is also present in the subjective beliefs data.

Figure 2. Willingness to Pay and Subjective Expected Payoffs



Notes: This figure is a binned scatter plot of willingness to pay (p) and subjective expected payoffs (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.

$\beta < 0$, which implies there should be a downward bias in the estimated risk-return relation. If this bias is strong enough, it can flip the sign of the relationship from positive to negative. Our data reveal that the bias is substantial: Column 1 of Table 3 shows that there is a negative relationship between expected return and risk. Figure 3 illustrates the negative correlation.[11] This result may initially appear puzzling in light of our earlier finding that subjects exhibit a lower valuation for the asset as perceived risk increases (Column 2 of Table 2). Our conceptual framework resolves the puzzle by showing that the negative relationship is caused by an omitted variable problem, which derives from the weak transmission of beliefs to WTP.

It is worth spelling out the intuition in more detail for how the econometrician can detect a negative risk-return relationship among a population of risk averse subjects. The reason is that when a risk averse subject's beliefs are not fully transmitted into her WTP,

[11] Appendix D shows that the risk-return relationship remains qualitatively unchanged when using the subjective probability for the lowest possible payoff ($60) as the perceived risk measure.

the *measured* subjective expected return $r_t = d_t - p_t$ no longer depends only on the risk premium; it also depends on the payoff expectation, as illustrated in equation (4). In other words, the subjective expected return based on reported beliefs is no longer equal to the more economically meaningful discount rate that represents risk compensation. Importantly, we can restore the positive risk-return relation by controlling for payoff expectations in a regression of expected return on risk. Column 2 of Table 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 sum, our results are consistent with Prediction 1.

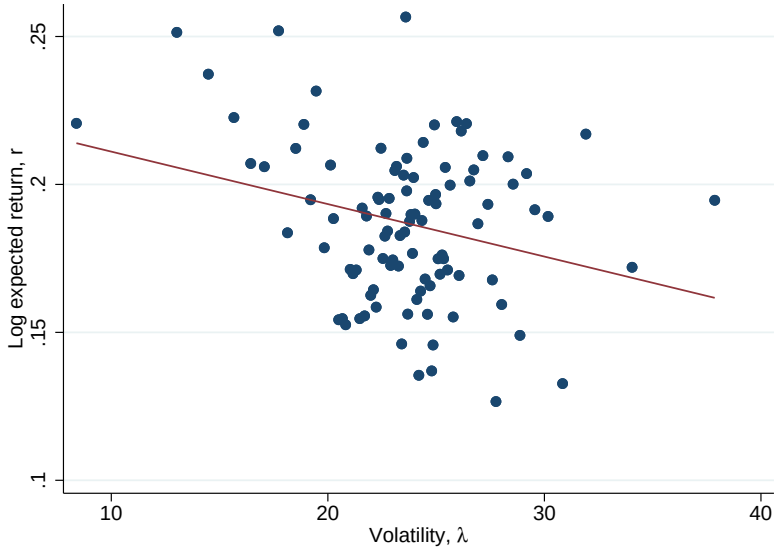
Table 3
Subjective Expected Returns, Subjective Expected Payoffs, and Perceived Risk

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

Notes: This table presents results from mixed effects regressions of subjective expected returns (r) on subjective expected payoffs (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 have thus shown that even when subjects are risk averse, the weak transmission of beliefs to WTP can create the illusion that risk premia implied by subjective expected returns are declining in risk. This phenomenon arises because subjective expected returns no longer represent only the discount rate which rises in bad times, but also the payoff expectation which declines in bad times. If the latter force dominates – which can occur with sufficient cognitive noise – the subjective expected returns will become procyclical, which is wholly consistent with our data.

Figure 3. Subjective Expected Returns and Perceived Risk



Notes: This figure is a binned scatter plot of the subjective expected return (r) and perceived risk (λ) controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300.

3.2.3 Subjective vs. objective expected returns

So far we have focused on analyzing the link between subjective expected returns and perceived risk. What do our results imply for the link between *objective* expected returns and risk? What would an econometrician who observes WTP and expectations from realized payoffs infer from our data? To investigate these questions we first define the objective expected return as

$$\begin{aligned} r_t^b &= d_t^b - p_t \\ &= \underbrace{(d_t - p_t)}_{r_t} - (d_t - d_t^b). \end{aligned} \tag{10}$$

where d^b is the log Bayesian expectation of the payoff. An econometrician can compute d^b, like we did, using the sequence of payoff realizations. Combining the subject's WTP with the Bayesian payoff expectation gives us r^b. The Bayesian expected return is conceptually similar to measures of objective expected returns in asset pricing computed using predictive regressions or Bayesian estimation.

Equation (10) shows that the difference between the objective and subjective expected returns, r and r^b, is driven by the departures of subjective payoff expectations from the Bayesian benchmark, namely, the difference between d and d^b. This immediately implies that the difference in the sensitivity of the objective and subjective expected return to risk depends on the relation between $d - d^b$ and risk. In Appendix Figure A.2, we show that $d - d^b$ is steeply declining in perceived risk with a coefficient of -0.0056 ($p < 0.01$). While the departure is centered around zero, subjects are too pessimistic in times of high perceived volatility and too optimistic in times of low perceived volatility. This suggests an overall overreaction in expectation formation.

The strong relationship between $d - d^b$ and risk also suggests that objective expected returns have a higher sensitivity to perceived risk compared to their subjective counterparts. The reason is that objective expected returns can be rewritten as subjective expected returns minus $d - d^b$ (see Equation (10)). Since the term $d - d^b$ strongly varies with risk, objective expected returns are especially sensitive to risk. Indeed, the coefficient of objective expected returns on risk is 0.005 ($p < 0.01$), which is larger in magnitude than the coefficient of subjective expected returns on risk shown in Column 1 of Table 3. Perhaps more strikingly, the coefficient for objective expected returns is positive (see Appendix Figure A.3 for a visualization). Therefore, an econometrician using predicted payoffs to form objective expected returns would identify a completely different risk-return relation compared to the one using elicited subjective expected returns and risk.

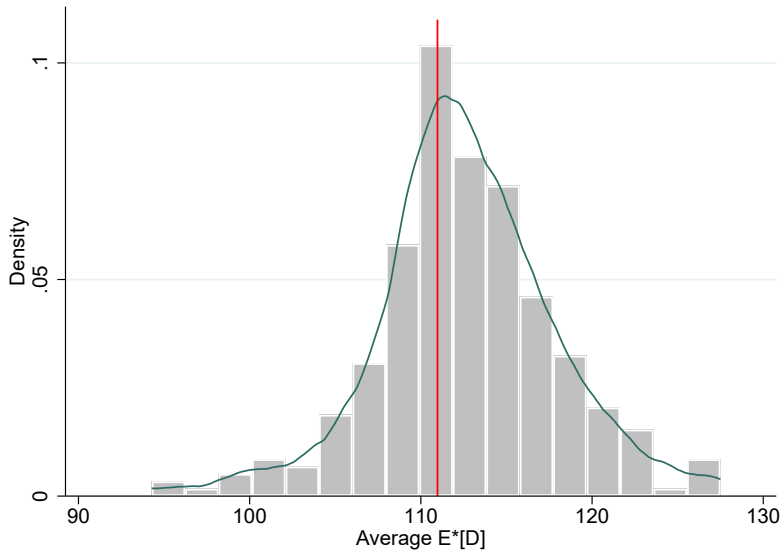
The intuition behind this stark contrast can be summarized as follows. In our experiment, when subjects perceive high volatility, they generally also expect a low payoff (one can think of this as subjects perceiving the state to be bad). In these situations, subjects lower their WTP both due to lower expected payoffs and higher discounts (Table 2). These effects are jointly stronger than the decline in the objective payoff expectation. Another way to interpret this result is that WTP fluctuates more strongly than Bayesian payoff expectations, despite the weak transmission of beliefs to WTP. As a result, the objective expected return increases with risk. These experimental results are qualitatively consistent with Nagel and Xu (2022b), who document countercyclical objective expected returns across several asset classes, but a lack of cyclicality in subjective expected returns.

To summarize, our controlled experimental setting generates data in which objective expected returns are higher and subjective expected returns are lower in bad times. The behavior of subjective expected returns is grounded in the weak transmission of beliefs to WTP. The differential sensitivity of objective vs. subjective expected returns to risk is driven by the systematic departures of subjective expectations from the Bayesian benchmark.

3.2.4 Objectivity of beliefs and the transmission of beliefs to willingness to pay

Earlier we highlighted that one advantage of our conceptual framework is that all empirical tests can be conducted regardless of whether beliefs are rational, where rational beliefs are defined as the beliefs of a Bayesian learner. Here, we test whether subjects whose subjective beliefs are closer to the objective Bayesian beliefs also act on their beliefs more aggressively.

Figure 4. Subject-level Average Expected Payoff



Notes: This figure plots the histogram and kernel density of the average expected payoff $\mathbb{E}^*[D]$ at the subject-level. Each observation is the average expected payoff of a subject across the eight elicitations. The vertical red line corresponds to the average objective (Bayesian) expectation. The sample size is 300.

Referring back to our summary statistics in Table 1, one salient fact is that the average deviation of subjective payoff expectations from Bayesian expectations is small. This can

be seen more clearly in Figure 4, which plots a histogram and kernel density of subject-level average expectations. The density is roughly centered at the true mean, but there is substantial heterogeneity. A large number of subjects are pessimists who expect payoffs below the objective expectation; there are also many subjects who are optimists and expect payoffs above the objective expectation.

We now examine, for each subject, whether the transmission of beliefs to WTP depends on how far their beliefs are from the objective benchmark. We conjecture that subjects whose beliefs are better calibrated will also be more confident in their beliefs. Given past work demonstrating that more confident investors exhibit a stronger sensitivity to beliefs, we hypothesize that subjects with better calibrated beliefs are more sensitive to these beliefs in their pricing (Giglio et al. (2021a)). We emphasize that precise control of the dividend process is key to conducting a test of this hypothesis. It would be difficult to implement such a test in the field where neither the investor nor the econometrician has access to the true data generating process of dividends.

To begin, 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 error$_s = \sum_{t=1}^{8} |d_{st} - d_t^b|$ where d_s is the subjective payoff for subject s and d^b is the Bayesian expectation. The median calibration error across all 300 subjects is 0.552. We define those subjects who are below the median – and thus exhibit beliefs that are relatively close to the objective benchmark – as *calibrated*. Those subjects who are above the median are defined as *miscalibrated*.

We then estimate the sensitivity of WTP to beliefs for each of the two subsamples. The first column of Table 4 shows that for the *calibrated* subsample, the coefficient of WTP on expected payoff is almost exactly 1. By contrast, in the second column we see that the coefficient magnitude is cut in half for the *miscalibrated* subsample. The third column shows that the difference is indeed statistically significant at the 1% level. Thus, our results are consistent with the conjecture that subjects who form beliefs that are closer to the objective benchmark also transmit these beliefs to their WTP in a stronger fashion. We emphasize that the results in Table 4 only provide correlational evidence that beliefs closer to the objective benchmark are transmitted more strongly into prices. In the next section, we discuss an

Table 4
Transmission of Beliefs for Calibrated and Miscalibrated Subjects

Sample:	Calibrated	Miscalibrated	Difference
p	(1)	(2)	(1) - (2)
d	1.068***	0.506***	0.562***
	(0.075)	(0.054)	(0.092)
Observations	1,200	1,200	2,400

Notes: This table presents results from mixed effects regressions of willingness to pay (p) on subjective expected payoffs (d) separately for *calibrated* (first column) and *miscalibrated* (second column) subjects. The difference in coefficients is reported in the third column. These regressions include a random effect for d 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.

experiment that provides a causal test.

4 Experiment 2

In this section, we exogenously manipulate cognitive noise and measure the impact on the transmission of beliefs to WTP. In particular, we test whether we can increase the transmission strength to the extent that the measured risk-return relationship becomes positive.

4.1 Experimental design

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

In order to test this hypothesis, we conduct a second experiment in which we endow subjects with objective beliefs about the next period dividend. Subjects do not need to learn because we explicitly provide them with the true payoff distribution. Our manipula-

tion 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.1.

The design allow us to test Prediction 2: reducing cognitive noise will strengthen the transmission of beliefs to WTP and, as a result, increase the coefficient on risk in a regression of expected returns on risk. Importantly, any change we detect across experiments in the estimated risk-return relationship must be due to a difference in the strength of transmission of beliefs to WTP. To see this, recall that the sensitivity of expected return to risk is given by $(1-x)\beta - (1-x)\gamma + \gamma$ (from equation 6), and the only value in this expression that varies across our two experiments is x. Why? The value of β, which captures the correlation between expected payoff and volatility, does not vary across experiments because we impose the beliefs that subjects report from Experiment 1 on those subjects in Experiment 2. The value of γ also remains unchanged across experiments: γ represents the risk aversion of our subjects, and we randomly draw subjects from the same population in each experiment. Taken together, any change in the estimated risk-return relationship across experiments can

be attributed to the reduction in cognitive noise.

As in Experiment 1, subjects in Experiment 2 are incentivized using the BDM mechanism and one of the questions is randomly chosen to be paid. Note that subjects in Experiment 2 therefore answer only 8 questions (compared with 16 in Experiment 1). To keep the incentives per question constant 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, and we want to hold incentives constant across experiments.

As in Experiment 1, we recruit subjects from Prolific and pre-registered the experiment on Aspredicted.org.[12] Subjects received $2 for completing the experiment, in addition 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.

4.2 Experimental results

Experiment 2 delivers a panel dataset with 2,400 observations that consist of eight distinct WTPs, elicited from each of the 300 subjects. Table 5 reports summary statistics from the dataset. Because we endow subjects with beliefs, we focus here only on the elicited WTP and the implied expected returns. We observe similar average WTP and expected returns across the two experiments, though both variables exhibit slightly greater dispersion in Experiment 2 compared to Experiment 1.

Table 5
Summary Statistics for Experiment 2

		Mean	p25	p50	p75	SD	Min	Max	N
Willingness to pay	P	97.00	79.80	99.90	115.00	23.22	60.00	150.00	2,400
Expected return	$\mathbb{E}^*[R] = \mathbb{E}^*[D]/P$	1.22	1.00	1.14	1.40	0.30	0.48	2.36	2,400

Notes: This table presents summary statistics for the variables in Experiment 2 that are different from Experiment 1. The sample consists of 300 subjects, each elicited for 8 payoff distributions, yielding 2,400 observations. P is the subject's reported willingness to pay for next period's dividend. $\mathbb{E}^*[R]$ is the expected return using the mean of the dividend distribution and the subject's willingness to pay.

We first test whether our experimental manipulation indeed strengthens the transmission

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

of beliefs to WTP. Column 1 in Table 6 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 our Prediction 2, the transmission of beliefs to WTP is causally strengthened when we endow subjects with beliefs, thereby eliminating the need for subjects to learn in order to arrive at a belief. 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 intercept of the regression is lower in Experiment 2 because the higher loading on payoff expectations captures a larger fraction of the level of the WTP.

Table 6

Willingness to Pay, Expected Payoffs, and Perceived Risk in both Experiments

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)
Observations	4,800	4,800

Notes: This table presents results from mixed effects regressions of willingness to pay (p) on expected payoffs (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.

The stronger transmission of payoff expectations to WTP remains significant after controlling for risk as shown in Column 2 of Table 6. The coefficient on d significantly increases from 0.610 in Experiment 1 to 0.817 in Experiment 2. These results suggest that cognitive noise from expectation formation explains about half of the underreaction of WTP to payoff expectations.

In line with Prediction 2, decreasing cognitive noise also increases the reaction of WTP to risk perceptions. Column 2 of Table 6 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, 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).

It is important to note that the transmission from beliefs to actions is still not as strong as predicted by the frictionless benchmark; the coefficient on d remains significantly below one in Experiment 2. Endowing subjects with objective beliefs clearly reduces cognitive noise, but it is not fully eliminated. In the discussion section, we speculate on other sources of cognitive noise that remain in our Experiment 2.

Table 7

Expected Returns, Expected Payoffs, and Perceived Risk in both Experiments

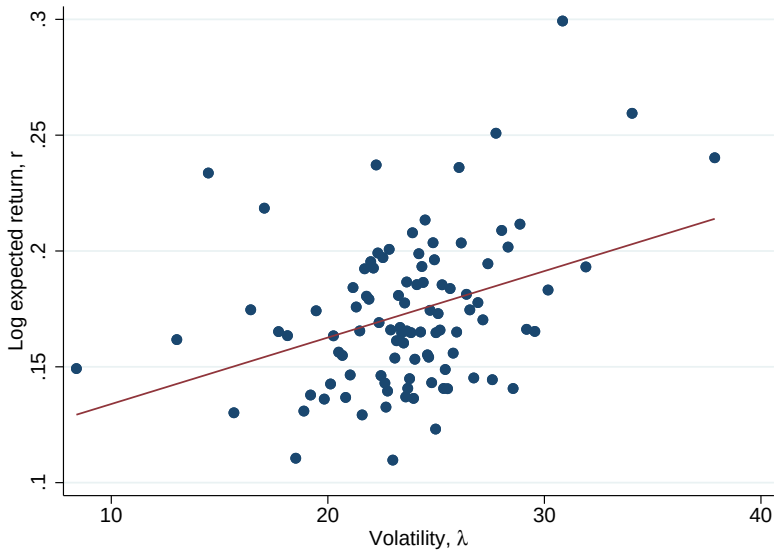
r	(1)	(2)
d		0.390***
		(0.050)
d x Exp2		−0.207***
		(0.066)
λ	−0.174**	0.188***
	(0.078)	(0.073)
λ x Exp2	0.418***	0.221*
	(0.121)	(0.121)
Exp2	−0.113***	0.910***
	(0.032)	(0.322)
Observations	4,800	4,800

Notes: This table presents results from mixed effects regressions of expected returns (r) on expected payoffs (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.

We now turn to assessing the causal effect of cognitive noise on the risk-return relationship. Equation (6) shows that an increase in x mitigates the omitted variable bias in the univariate risk-return relation. Mitigating the omitted variable bias will increase the slope of the risk-return relation, and it can potentially restore a positive risk-return relation. Here

we report results from estimating the risk-return relationship, which by design uses the same set of beliefs from Experiment 1. Column 1 of Table 7 shows that the loading on λ becomes positive in Experiment 2, where subjects price the asset based on the objective payoff distribution. The fact that the sign of the relationship flips demonstrates that endowing subjects with objective beliefs – rather than having them form subjective expectations as in Experiment 1 – is associated with a substantial reduction in cognitive noise. We are able to make this inference because β and γ are held constant across experiments, and thus all of the action must be through an increase in x, which we interpret as a reduction in cognitive noise. Figure 5 demonstrates the positive risk return relationship. It is worth noting that the relationship looks markedly different compared with the risk-return relation from Experiment 1 (displayed in Figure 3).

Figure 5. Expected Returns and Risk in Experiment 2



Notes: This figure is a binned scatter plot of expected returns (r) and risk (λ) in Experiment 2 controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300.

Finally, recall that cognitive noise induces a positive correlation between expected payoffs and expected returns. Since we predict less cognitive noise in Experiment 2, the correlation between expected payoff and expected return should be attenuated towards zero. We find

that this is indeed the case. In Column 2 of Table 7, we regress expected return on both expected payoff and risk, and find that the coefficient on expected payoff is significantly closer to zero in Experiment 2 compared to Experiment 1. In other words, the impact of payoff expectations on expected returns induced by the weak transmission becomes smaller in Experiment 2, which is consistent with equation (4).

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. The increased responsiveness is so large that it flips the sign of the risk-return relationship. Our results suggest that subjects rely much less on a cognitive default parameter when they are endowed with objective beliefs.

5 Discussion

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

5.1 Implications for asset pricing models

Over the past few years several models have been proposed to quantitatively match both asset prices and survey expectations (e.g., Barberis et al. (2015), Hirshleifer et al. (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. Our focus in this paper takes expectations as given, and examines how these expectations propagate into actions. One key finding that emerges from our experiments is that the transmission of beliefs to actions is far from 1-to-1. We believe that this result should motivate future theoretical work to explicitly incorporate weak transmission into the investor's decision process, in order to potentially improve quantitative

fits to the data.[13]

An obvious concern, however, is that injecting deviations from rationality in both expectation formation and expectation transmission gives the modeler too much flexibility. Indeed, the number of non-rational expectations models is large enough, and adding an additional 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 4 we show that subjects who state beliefs closer to the rational benchmark are the same subjects whose valuations are more sensitive to these beliefs. While more empirical data is needed, our results suggest that the belief formation process can partially constrain the degree of weak transmission.[14]

Our results also provide guidance for asset pricing models with 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. This imperfect information approach is common in the asset pricing literature. For instance, several models generate realistic asset price dynamics by relying on the fluctuations in beliefs that result from Bayesian learning about the stochastic state (e.g., Veronesi (1999), Johannes et al. (2016), Ghaderi et al. (2022)). Our results suggest that learning models in particular may benefit from incorporating the weak transmission of beliefs to actions.

More broadly, our results suggest that the degree of weak transmission depends on whether investors have state uncertainty. Recall that in Experiment 1, subjects must learn from the payoff process, while in Experiment 2, we shut down the learning channel and endow subjects with beliefs. We find that when subjects are endowed with beliefs, their asset valuations are substantially more sensitive to the endowed beliefs. Because the beliefs are held constant across experiments, our results provide evidence that belief formation causally affects how investors act on those beliefs.

[13]For related work in macroeconomics, see Khaw et al. (2017) for a theoretical model that incorporates inattentive adjustment of actions.

[14]See also Andries et al. (2022), who show that experimental subjects are more likely to underreact in their investment decisions when forming extrapolative expectations compared to rational expectations.

Applying this intuition to the field, our Experiment 1 corresponds to a real-world environment in which investors form expectations about the stock market. For instance, at the beginning of the COVID-19 crisis in March 2020, investors formed a subjective assessment of how the stock market would behave going forward. Giglio et al. (2021b) show that investors updated their beliefs and became more pessimistic about the stock market in the short term. However, these investors did not adjust their portfolios nearly as much as frictionless asset pricing models would predict. This finding is consistent with the weak transmission of beliefs to actions, documented in our Experiment 1. In contrast, Experiment 2 constructs a counterfactual to Experiment 1 that is impossible to implement in the field. It allows us to test how investors would have responded if they were endowed with the objective payoff distribution. This is akin to asking how investors would have behaved in the COVID-19 crisis, if they were given an objective probability distribution of the short-run stock market performance. The results from Experiment 2 suggest that investors would have reacted more strongly to these (objective) beliefs and would have adjusted their portfolios more aggressively.

Our paper is also related to a recent strand of literature that quantifies a demand system using portfolio holdings of different investors. One key finding from this literature is that investors' asset demand is very inelastic compared to predictions from standard asset pricing models (Koijen and Yogo (2019)). The weak transmission due to cognitive noise, which is the foundation for the biased risk-return relationship in our work, may be helpful in explaining the inelasticity. Specifically, cognitive noise can provide a potential microfoundation for the observed low elasticity of demand for assets in response to expected return fluctuations (e.g., Haddad et al. (2021), Chaudhry (2022)).

5.2 Connections with the survey literature

Our results directly connect to the growing literature on survey data in finance. Recent work has examined the specific factors that investors say are important to their personal investment decisions (Bender et al. (2021), Chinco et al. (2021)). For example, beyond expected return and risk, Choi and Robertson (2020) find that investors report other factors such as time left until retirement as important determinants of their portfolio decisions. While we focus our analysis on the sensitivity to expected payoff and perceived risk, it is

possible that investors may not act as aggressively on other factors that they state in surveys. Indeed, if cognitive noise is in part responsible for affecting an investor's valuation of an asset conditional on risk and return, then it seems plausible that cognitive noise would also be present when investors use additional factors to arrive at asset valuations.

The survey conducted by Giglio et al. (2021a) is more closely related to our specific experimental paradigm. Two salient, but puzzling facts that Giglio et al. (2021a) document are that (i) portfolio equity shares do not respond as strongly to reported beliefs as predicted by frictionless asset pricing models, and (ii) investors expect low stock returns when they perceive aggregate risk to be high. Our experimental results may help connect the two facts and provide a potential explanation for them. Suppose that investors increase their perception of risk, and they do not fully adjust their willingness to pay downward because of weak transmission. Investors are then less likely to sell equity compared to the case in which they change their willingness to pay in a 1-1 fashion with beliefs. It follows that equity prices will remain relatively higher in equilibrium, compared to the case in which there is stronger transmission of beliefs to actions. In such a scenario, expectations of the future level of the stock market (P_{t+1}) are low, while the current price level of the stock market (P_t) remains relatively high. Since subjective return expectations are simply $\mathbb{E}_t^*[P_{t+1} + D_{t+1}]/P_t$, these two forces result in low return expectations. In sum, the weak transmission of beliefs to actions provides one potential explanation for why investors report low expected returns when they perceive aggregate risk to be high (Giglio et al. (2021a)), and why subjective return expectations lack the cyclicality of their objective counterparts (Nagel and Xu (2022b)).

In the same dataset, Giglio et al. (2021a) document a positive relation between GDP growth expectations and subjective return expectations. This fact presents a challenge for aggregate asset pricing models that generate high volatility in valuation ratios (e.g., the price-dividend ratio of equity). The reason is that "bad times" in models are typically characterized by low fundamental growth expectations (i.e., low GDP growth expectations) and high discount rates (i.e., high expected returns), both of which lower prices. Similar to the evidence in Giglio et al. (2021a), we find a positive relation between subjective expected returns and payoff expectations (see Appendix Figure A.4). Naturally, our experimental asset is simpler than aggregate equity, but coupled with our conceptual framework, we can resolve

the puzzle in our setting. Intuitively, the weak transmission induces a positive correlation between expected payoffs and subjective expected returns because WTP does not increase 1-to-1 with expected payoffs. Economically, this means that the subjective expected return no longer represents the discount rate used for valuations, since subjective expected returns are confounded by the payoff expectation.

Our results also connect with research on analyst forecast surveys. De La O and Myers (2021) show that analyst forecasts of cash flows explain almost all the variation in stock market valuations while return expectations play a much smaller role. In Appendix C, we show that payoff expectations explain a higher fraction of price variation under subjective beliefs compared with Bayesian beliefs in our experiment. Such a fact is qualitatively consistent with the pattern documented by De La O and Myers (2021) in the field.

5.3 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 heterogeneous 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 some of the cognitive noise arises from uncertainty about expectations, perhaps because subjects have difficulty implementing Bayes' rule (Kuhnen (2015)), Ben-David et al. (2019)). But importantly, in Experiment 2, we show that shutting down uncertainty about beliefs still leads to weak transmission. We speculate that the noise in Experiment 2 arises primarily from integrating beliefs about payoffs with perceived risk to arrive at a valuation.

5.4 Limitations

We have argued that the one period nature of the asset in our experiment 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 experiment is because the degree of extrapolation, and more generally, expectational errors, may depend on the horizon of the forecast (Giglio and Kelly (2018), Da et al. (2021), De Silva and Thesmar (2021)). 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. (2021), 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.

6 Conclusion

Survey data on subjective beliefs have recently opened up a vibrant area of research in asset pricing (Adam and Nagel (2022)). 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 exploring the implications of investors who do not fully act on their stated beliefs. We show theoretically that the weak transmission of beliefs to actions induces a substantial bias in even the most basic asset pricing tests.

Our experimental data provide strong support for the prediction of a downward bias in the risk-return relation. Subjects in our experiment are indeed insensitive investors, and we find that expected returns systematically decline in perceived risk – despite the fact that the average subject is risk averse. Our framework also provides a recipe for restoring the positive risk-return relation: include expected payoff in the regression of expected returns on

risk. In our data, adding this control flips the sign of the risk-return relation from negative to positive.

Because our experiment shuts down institutional frictions by design, we identify the source of the weak transmission as a psychological friction. In particular, we interpret the weak sensitivity as arising from cognitive noise, and we find strong evidence that cognitive noise causally affects the risk-return relation. We interpret the noise 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.

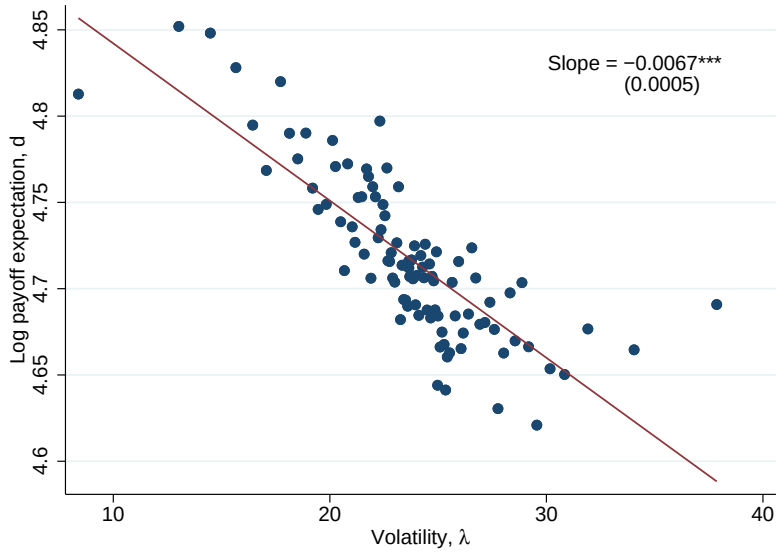
The fact that we document the weak transmission of beliefs to actions in a controlled experimental setting points to the idea that weak transmission may be a fundamental component of the investor's decision process. As such, one natural path for future work is to incorporate this ingredient into existing asset pricing models and assess the change in quantitative fit. Another path forward is to better understand the importance that institutional frictions play in generating weak transmission in the field. Because institutional frictions are shut down in our experiment, the bias may be even more substantial in tests using data from the field.

Appendix

A Additional figures

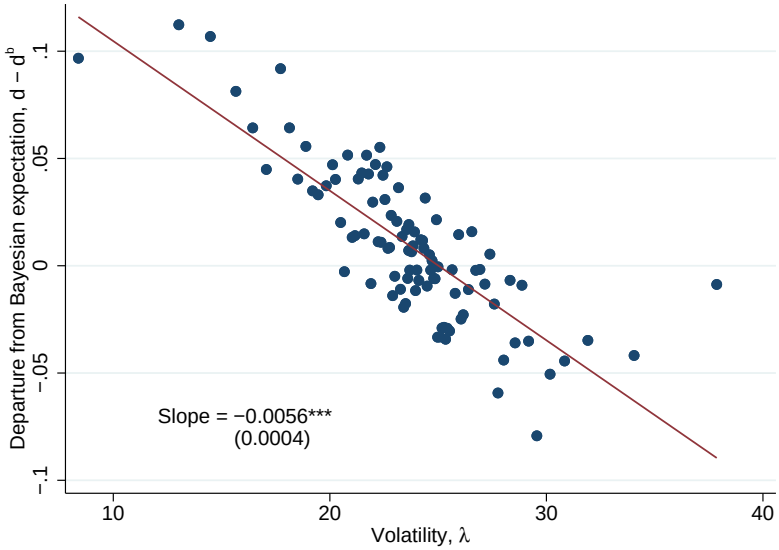
Here, we provide several additional figures that are referenced in the main body of the paper.

Figure A.1. Subjective Expected Payoffs and Perceived Risk



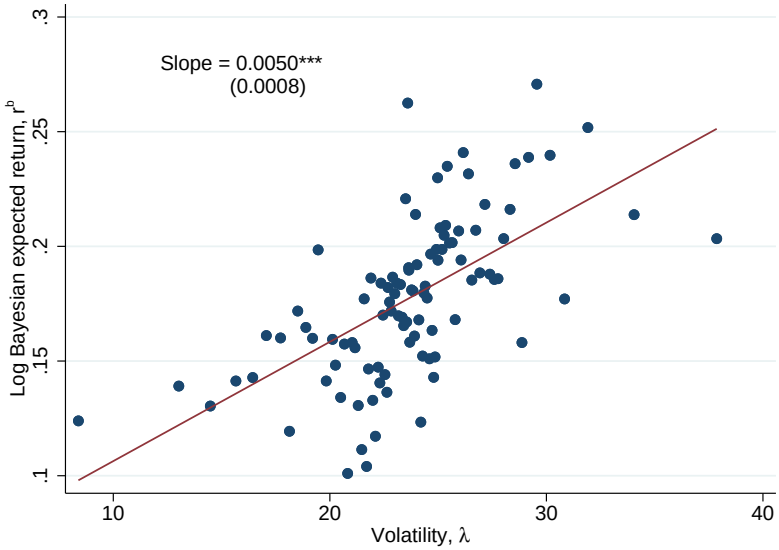
Notes: This figure is a binned scatter plot of subjective expected payoffs (d) and perceived volatility (λ) controlling for subject fixed effects. The reported slope results from a mixed effects regression of d on λ. The regression includes a random effect for λ as well as for the intercept. The standard 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.

Figure A.2. Departures from Bayesian Expectations and Perceived Risk



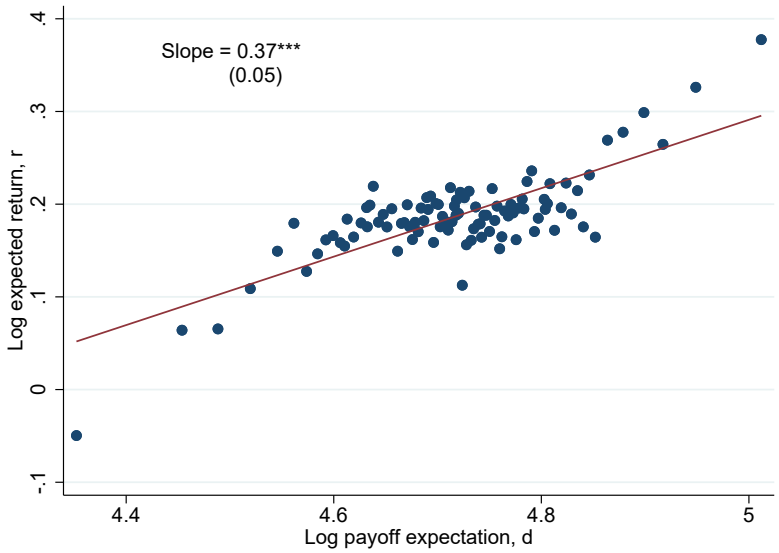
Notes: This figure is a binned scatter plot of departures of subjective expected payoffs from the Bayesian benchmark ($d - d^b$) and perceived risk (λ) controlling for subject fixed effects. The reported slope results from a mixed effects regression of $d - d^b$ on λ. The regression includes a random effect for λ as well as for the intercept. 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.

Figure A.3. Bayesian Expected Returns and Perceived Risk



Notes: This figure is a binned scatter plot of Bayesian expected returns (r^b) and perceived risk (λ) controlling for subject fixed effects. The reported slope results from a mixed effects regression of r^b on λ. The regression includes a random effect for λ as well as for the intercept. 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.

Figure A.4. Subjective Expected Returns and Subjective Expected Payoffs



Notes: This figure is a binned scatter plot of subjective expected returns (r) and subjective expected payoffs (d) controlling for subject fixed effects. The reported slope results from a mixed effects regression of r on d. The regression includes a random effect for d as well as for the intercept. 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.

B Derivations for the conceptual framework

B.1 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.2 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. Rearranging equation (5) 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 (3) to obtain

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

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 $\beta < 0$.

C Expected cash flow and return effects

We decompose the variation in WTP into expected payoff and expected return effects and report results in Table A.1. This exercise is similar to the Campbell and Shiller (1988) decomposition of the price-dividend ratio into expected dividend growth and expected return effects.

Table A.1
Decomposition of Variation in WTP

Bayesian		Subjective	
d^b	4%	d	25%
- r^b	96%	- r	75%

Notes: This table shows the decomposition of variation in willingness to pay p into expected payoff and expected return effects using the identities $p = d - r$ and $p = d^b - r^b$. The numbers represent $\frac{cov(\tilde{q},p)}{var(p)}$ where $\tilde{q}$ is one of d^b, d, - r^b, and -r. The sample size is 2,400 and the number of subjects is 300. The data are from Experiment 1.

D Results using tail risk

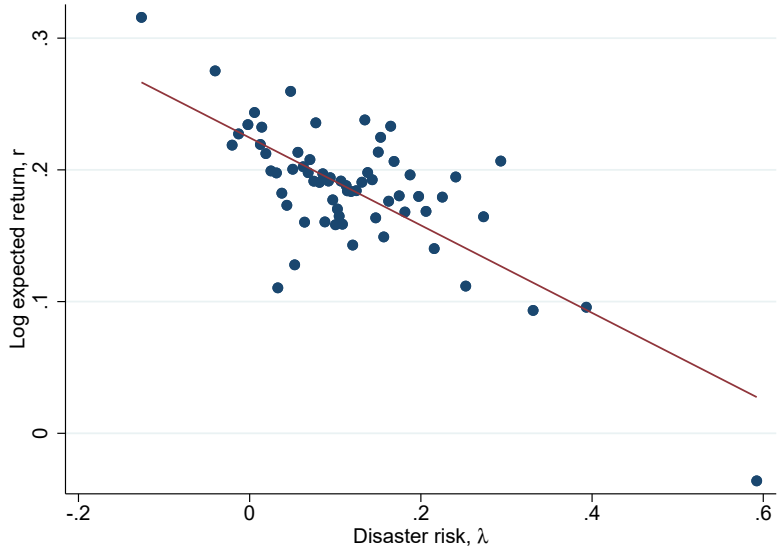
In this section, we present the risk-return relationship in Experiment 1 using the subjective probability of the lowest payoff ($60) as the measure of perceived risk. This is a measure of tail risk and is similar to the measure of disaster risk in Giglio et al. (2021a). Table A.2 and Figure A.5 show that the results are similar to our main results from Table 3, in which we use volatility as the perceived risk measure.

Table A.2
Subj. Expected Returns, Subj. Expected Payoffs, and Perceived Tail Risk

r	(1)	(2)
d		0.401^{***}
		(0.051)
λ	-0.210^{**}	0.186^{***}
	(0.047)	(0.060)
Observations	2,400	2,400

Notes: This table presents results from mixed effects regressions of subjective expected returns (r) on subjective expected payoffs (d) and the perceived probability of the lowest payoff (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 data are from Experiment 1. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Figure A.5. Subjective Expected Returns and Perceived Tail Risk



Notes: This figure is a binned scatter plot of subjective expected returns (r) and the perceived probability of the lowest payoff (λ) controlling for subject fixed effects. The sample size is 2,400 and the number of subjects is 300. The data are from Experiment 1.

E Heterogeneity across subjects

Here, we analyze the sources of variation in the data from Experiment 1 and report the fractions of variation explained by time and subject fixed effects in Table A.3. Consistent with the field evidence from Giglio et al. (2021a), fixed effects explain a large fraction of WTP, subjective return expectations, and perceived risk. Payoff expectations have a larger component driven by time rather than subject fixed effects. This result likely reflects the commonality in learning.

Table A.3
Decomposition of Variation: Subject and Time Fixed Effects

R^2 in %	Time FE (1)	Subject FE (2)	Both FE (3)
d	24.8	8.1	36.2
p	6.1	65.6	71.7
r	0.6	68.3	68.8
λ	3.6	54.7	58.3
Observations	2,400	2,400	2,400

Notes: This table reports the R^2s corresponding to the regressions of the variables displayed in the first column on time fixed effects, subject fixed effects, or both. The dependent variables are subjective expected payoffs (d), willingness to pay (p), subjective expected returns (r), and perceived volatility (λ). The data are from Experiment 1.

REFERENCES

Adam, Klaus, and Stefan Nagel, 2022, Expectations data in asset pricing, NBER Working paper.

Afrouzi, Hassan, Spencer Yongwook Kwon, Augustin Landier, Yueran Ma, and David Thesmar, 2021, Overreaction in expectations: Evidence and theory, SSRN Working paper.

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, 2022, 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, 2019, Expectations uncertainty and household economic behavior, NBER Working paper.

Bender, Svetlana, James J Choi, Danielle Dyson, and Adriana Z Robertson, 2021, Millionaires speak: What drives their personal investment decisions?, *Journal of Financial Economics* forthcoming.

Beutel, Johannes, and Michael Weber, 2022, Beliefs and portfolios: Causal evidence, SSRN Working paper.

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

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.

Campbell, John Y, and Robert J Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Chaudhry, Aditya, 2022, The causal impact of macroeconomic uncertainty on expected returns, SSRN Working paper.

Chinco, Alex, Samuel M Hartzmark, and Abigail B Sussman, 2021, A new test of risk factor relevance, *Journal of Finance* forthcoming.

Choi, James J, and Adriana Z Robertson, 2020, What matters to individual investors? Evidence from the horse's mouth, *Journal of Finance* 75, 1965–2020.

Cochrane, John H, 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

Cochrane, John H, 2017, Macro-finance, *Review of Finance* 21, 945–985.

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 La O, Ricardo, and Sean Myers, 2021, Subjective cash flow and discount rate expectations, *Journal of Finance* 76, 1339–1387.

De Silva, Tim, and David Thesmar, 2021, Noise in expectations: Evidence from analyst forecasts, NBER Working paper.

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, 2021, Cognitive uncertainty, NBER Working paper.

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.

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

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, 2021, 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.

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

Nagel, Stefan, and Zhengyang Xu, 2022b, Dynamics of subjective risk premia, NBER Working paper.

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 Mete Kilic

Additional Tables and Figures

Not for publication

IA.1 Screenshots of the experiments

IA.1.1 Experiment 1

Subjects had full information of the dividend distribution in both states. The distributions were displayed to subjects before the first dividend realization and in each elicitation period.

Figure IA.1. Distribution in the *good* state

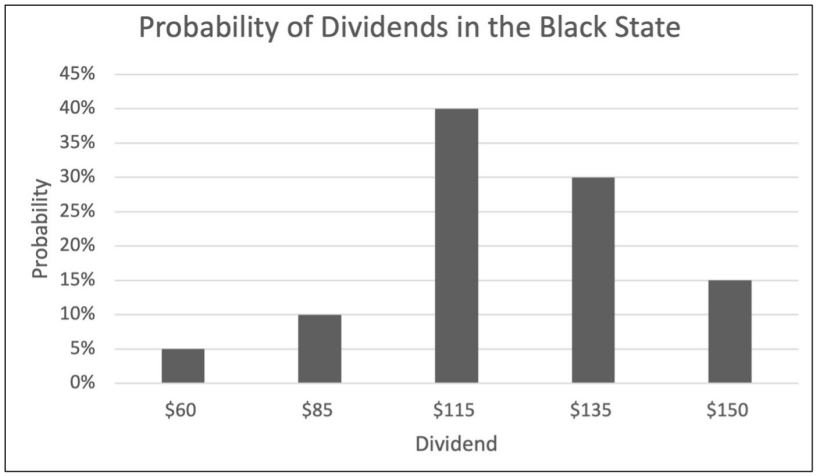
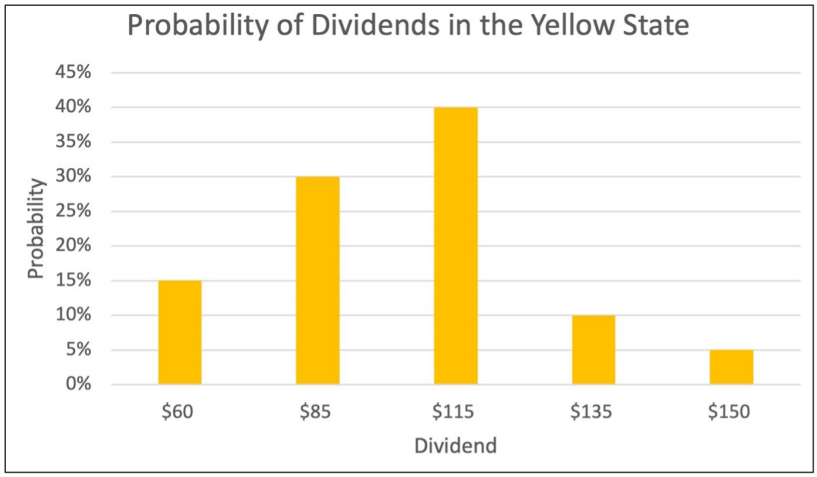
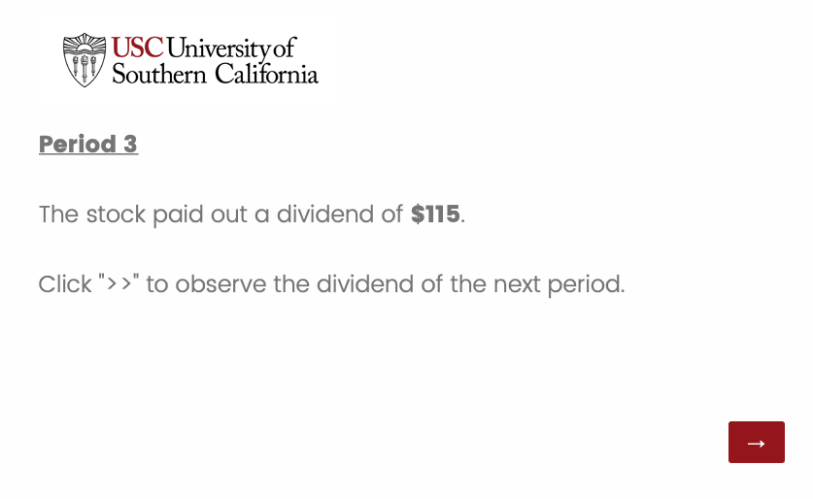


Figure IA.2. Distribution in the *bad* state



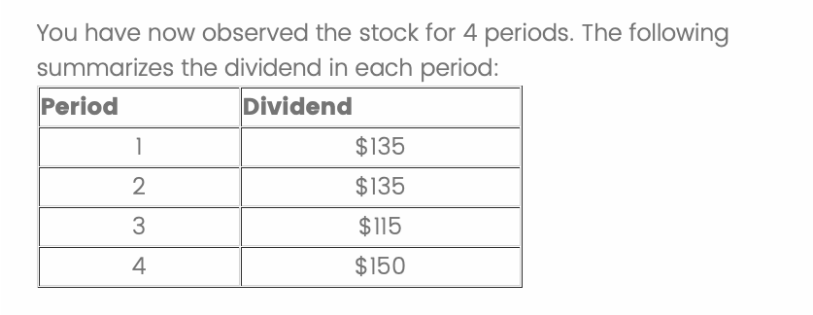
Below is a screenshot showing how dividend realizations were displayed to subjects in each period:

Figure IA.3. Dividend realization



After observing the dividend realizations over the course of several periods, subjects were asked to answer two questions. While answering these questions, they received an overview of the full history of dividend realizations:

Figure IA.4. Overview of history of dividend realizations



Subjects were able to report the probability that they attached to each dividend outcome. The ordering of the buckets (i.e., highest to lowest or lowest to highest) was randomized across subjects. The probability of each bucket was restricted to [0%, 100%] and the sum of the five probabilities was required to add up to 100%. Subjects were able to input their willingness to pay using a slider. This slider had to be initiated by the subject by clicking on the slider. The screenshots below show how the slider appeared before and after initiation:

Figure IA.5. Elicitation of payoff expectations

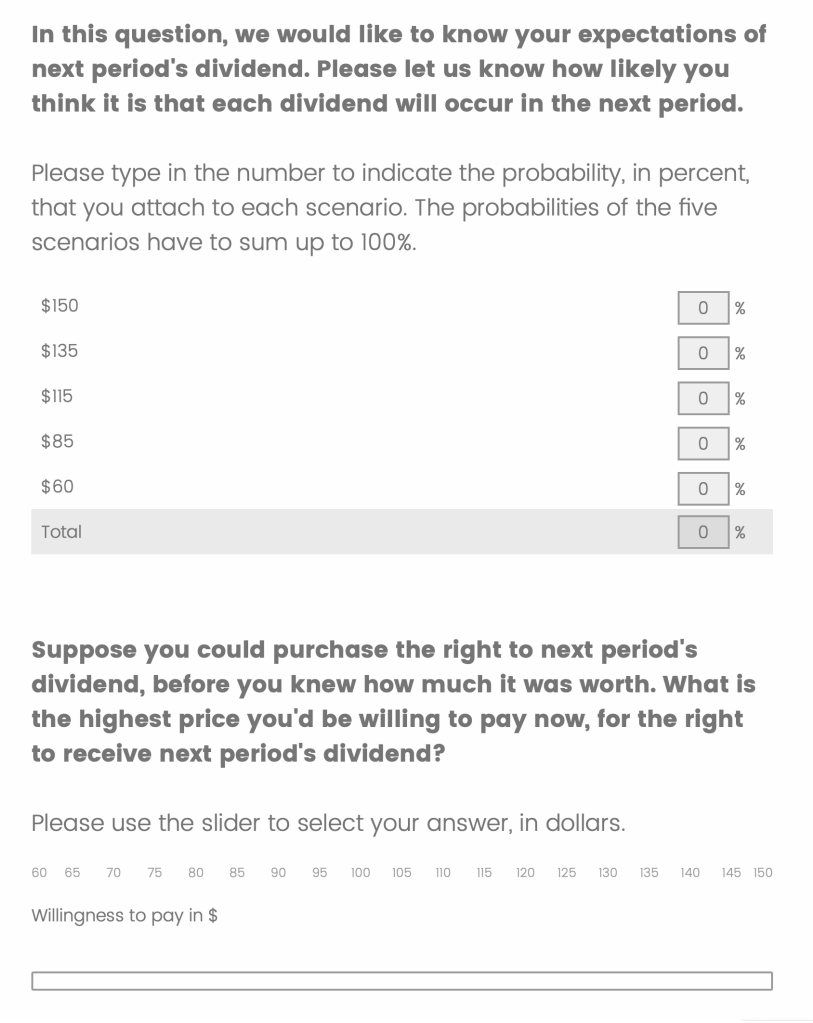
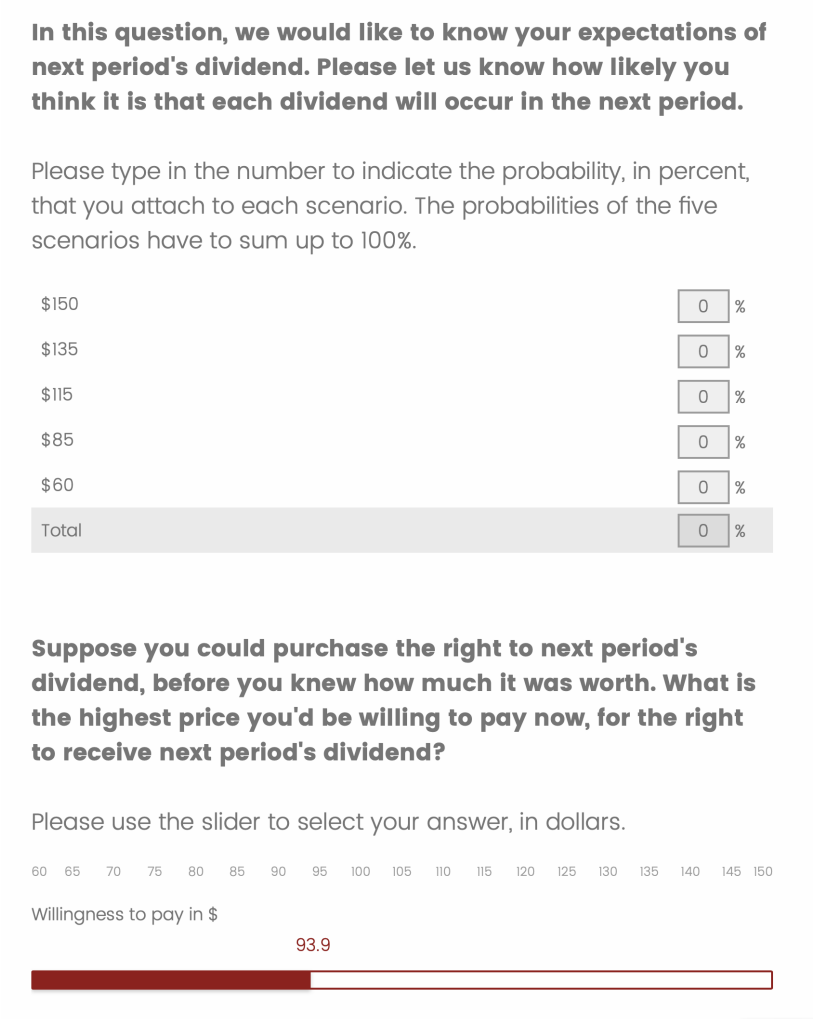


Figure IA.6. Elicitation of willingness to pay (after initiation)



IA.1.2 Experiment 2

Subjects were shown the probability of each dividend outcome. The ordering of the buckets (i.e., highest to lowest or lowest to highest) was randomized across subjects. Subjects were able to input their willingness to pay using a slider. This slider had to be initiated by the subject by clicking on the slider. The screenshots below show how the slider appeared before and after initiation:

Figure IA.7. Display of payoff distribution

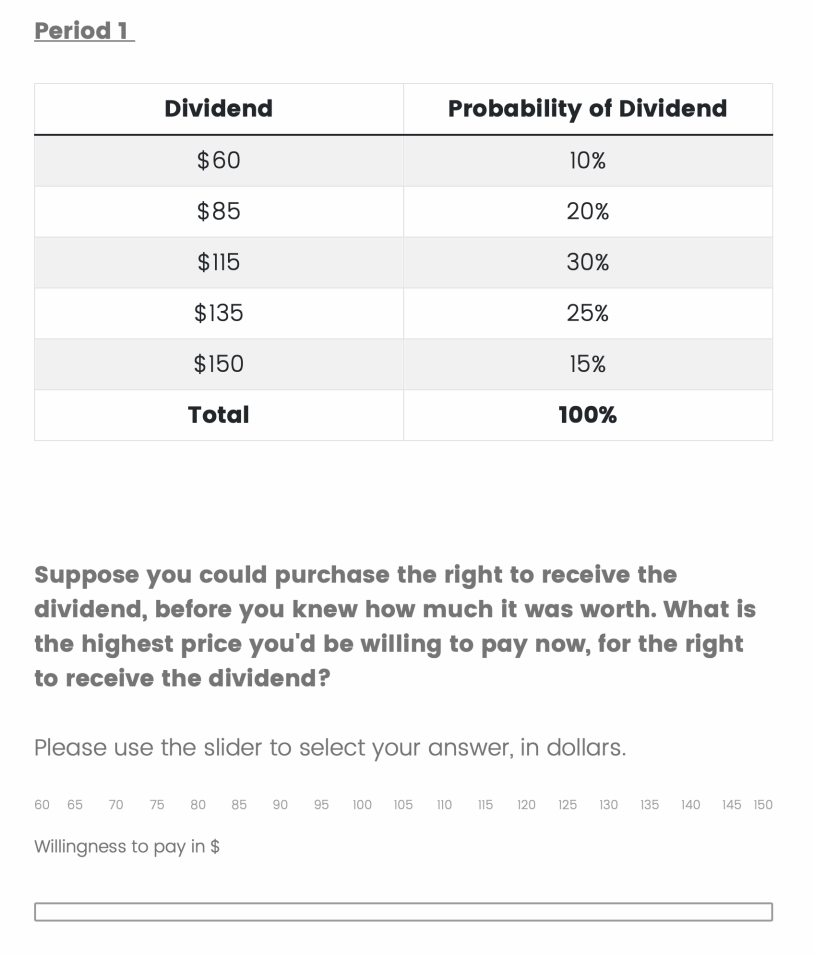
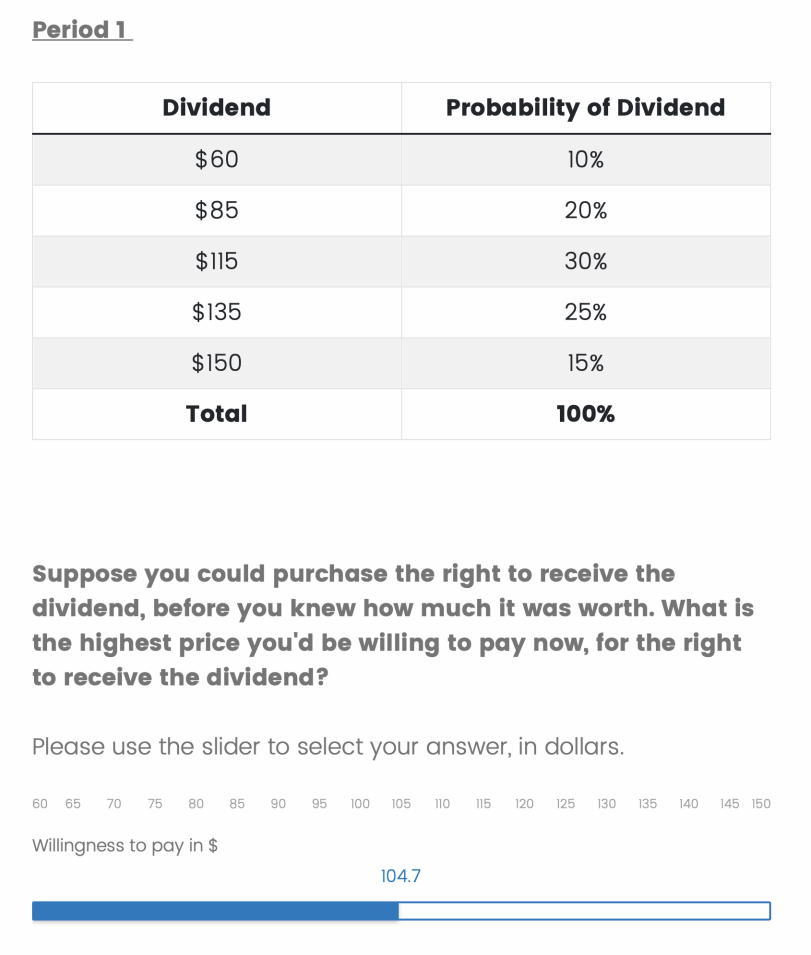


Figure IA.8. Elicitation of willingness to pay (after initiation)



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