Discounting Less in Bad Times:

Shining the Light on Cash Flow Expectations

Constantin Charles, Cary Frydman, and Mete Kilic[1]

November 2021

ABSTRACT: Research using survey data has found that respondents report lower expectations of future stock returns in bad times. This empirical pattern conflicts with the predictions of leading rational asset pricing models, where investors demand higher returns in bad times. We hypothesize that departures from rational cash-flow expectations can help reconcile the mismatch between theory and survey data. We test this hypothesis in an experiment that enables us to precisely control information about the cash flow process. Subjects are incentivized to report a time series of expected cash flows and asset valuations, which we use to infer discount rates. We find that discount rates and perceived risk are strongly negatively correlated; in contrast, a rational risk averse agent in our experiment should exhibit a strong positive relationship. We then document a new fact: when perceived risk is higher, subjects apply lower discount rates and are also willing to pay less for future cash flows. We argue that the mechanism which generates this new fact operates through distorted beliefs about cash flows. Overall, our results point to the importance of jointly modeling subjective expectations of cash flows and subjective expected returns.

[1] All three authors are from the University of Southern California, Marshall School of Business.
Charles: ccharles@marshall.usc.edu, Frydman: cfrydman@marshall.usc.edu, Kilic: mkilic@marshall.usc.edu. We are grateful to Ricardo De La O, David Hirshleifer, Lawrence Jin, and Selale Tuzel for helpful comments and discussions.

Over the past several years, there has been a surge of interest in using survey data to measure investor expectations about stock market returns (Brunnermeier et al., 2021; Adam, Mateveev, and Nagel, 2021). Greenwood and Shleifer (2014) synthesize data from multiple surveys and demonstrate that when past stock market returns have been high, investors and CFOs expect high future returns. Conversely, when past returns have been low, survey respondents expect low returns. This evidence clashes with a central prediction of prominent rational asset pricing models, in which investors demand (and expect) high returns when recent returns have been low. One path to resolve the mismatch between theory and data is to discount the survey responses, on the grounds that questions are non-incentivized and respondents are confused about expected returns. This view, however, is quickly being eroded by a new generation of survey data, which reinforces existing puzzles and reveals new facts.

Giglio et al. (2021a; 2021b) merge survey data on expectations with microdata on portfolio choice. With this new dataset, the authors reveal that survey expectations are highly informative about actual portfolio choice: those investors who expect higher stock market returns allocate a higher equity share in their portfolio. Giglio et al. (2021a) also provide a battery of new facts about the relationship between expected returns, perceived risk, and portfolio allocations. In line with the Greenwood and Shleifer (2014) evidence, Giglio et al. (2021a) show that when investors perceive high disaster risk, they expect *low* future returns – a relationship opposite to the prediction of rare disaster models. Additionally, portfolio allocations are too stable given the variation in stated expectations, cash flow growth expectations covary with expected returns, and there is substantial variation in expectations across investors.

In this paper we revisit several of these new facts in a controlled experimental setting, with an emphasis on understanding how discount rates vary with perceived risk. We design our experiment to provide three main advantages that complement data from past surveys. First, rather than assuming the expected returns reported by survey respondents are the same as discount rates, here we infer the discount rate in an incentive compatible manner. To do so, we elicit both a subject's willingness to pay for an asset and expectations of the asset's cash flows; these two ingredients enable us to back out the implied discount rate for each subject.

Second, we exogenously set the cash flow process and control the subject's information set; we can then make quantitative statements about how cash flow expectations and discount rates deviate from that of a risk averse Bayesian investor[2]. Finally, we incentivize subjects to price a one-period dividend strip in a partial equilibrium setting. This allows us to study the relationship between expectations and valuations at the subject level, without requiring the respondent to be the marginal investor. The one-period nature of the asset further simplifies the environment by shutting down the need for subjects to form cash flow expectations at longer horizons.

These design features give rise to a clean environment to investigate the basic building blocks of asset pricing. We ask subjects to report the full distribution of their beliefs about next period's dividend, and then we elicit their willingness to pay (or certainty equivalent) for the subjective cash flow distribution. The ratio of the subject's cash flow expectation to their willingness to pay is the discount rate, and we analyze how this quantity varies with a subject's perception of risk.

Our main results can be summarized as follows. Discount rates vary negatively with perceived risk, both within and across subjects. This result stands in stark contrast to the predictions of a rational benchmark model, in which a risk averse subject should discount cash

[2] See Afrouzi et al. (2021) for a recent experimental study that also exogenously sets the stochastic process that subjects need to forecast, in order to study deviations from Bayesian learning.

flows at a *higher* rate when perceived risk is high. The data are instead consistent with past surveys which find that when respondents perceive high risk, they expect low returns. We then show that willingness to pay and subjective cash flow expectations decline in perceived risk. Because the discount rate falls with perceived risk, it follows that cash flow expectations must be falling faster than willingness to pay (so that the ratio of the two also declines with perceived risk). Indeed, we find that as perceived risk becomes higher, subjects switch from overestimating to underestimating cash flows. This specific departure from rational expectations drives the observed negative relationship between discount rates and perceived risk.

In addition to matching the empirical relationship between subjective discount rates and risk from surveys, we also find that expected *realized* returns vary positively perceived risk. This pattern is observationally equivalent to the empirical data that motivates the leading dynamic asset pricing models, which induce variation in equity valuations using discount rate variation. Time variation in risk aversion in Campbell and Cochrane (1999), time varying volatility in Bansal and Yaron (2004), and time varying disaster risk in Gabaix (2012) and Wachter (2013) all generate high expected returns in times of high risk, contributing to the decline in valuations in bad times.

We show that the higher statistical expectation of realized returns in our experiment, e.g., implied by a predictive regression in the spirit of Campbell and Shiller (1988), does not imply that agents discount cash flows at a higher rate. Rather, we find that the low valuations in bad times are driven by excessively low cash flow expectations, and that discount rates are actually lower in times of high risk. The cash flow effect dominates the cyclicality of valuations leading to procyclical prices. Contrary to conventional wisdom in the asset pricing literature (e.g., Cochrane, 2011), we find that valuations are not driven by higher discounting in bad times. If anything, discount rate variation makes prices *less* volatile, all else equal. The lower discount rates in bad times can counteract the excessively pessimistic cash flow expectations, effectively serving as a cushion on which prices can fall. Conversely, the higher discount rates in good times will temper the high prices driven by excessively optimistic cash flow expectations[3].

Because we elicit the full distribution of cash flow beliefs from each subject over time, we have rich data to analyze departures from rational expectations. Interestingly, the deviations from rational cash flow expectations that we document are not easily detectable when simply looking at the subjective distribution of beliefs: on average, subjects report a distribution of cash flows that is reasonably close to the Bayesian benchmark. The systematic nature of the irrational expectations is revealed only when disaggregating the data across subjects and over time. When risk is high, subjects are overly pessimistic about cash flows, and when risk is low, subjects are overly optimistic. The subjective expectations are largely driven by wrong beliefs about the probability of receiving the lowest possible cash flow, which has an objectively low probability. When risk is high, subjects overestimate the probability of this low cash flow by 5.8 percentage points (relative to a rational benchmark of 10.4%), but when risk is low, subjects underestimate the probability of this cash flow by 3.2 percentage points (relative to a rational benchmark of 9.6%).

Our paper contributes to two different strands of literature. First, our results build directly on a recent set of papers in finance and macroeconomics that has returned to surveying investors and professional forecasters about their expectations (Greenwood and Shleifer, 2014; Kuchler

[3] In our experiment, we define good (bad) times as times of low (high) risk for the one-period asset. We note that while the asset pricing models discussed above induce risk premia through systematic risk in the economy, there is no notion of systematic risk in our experiment. Therefore, our statements regarding the attitude of a rational Bayesian agent towards risk implicitly assume that subjects perceive the experimental asset's risk as being non-diversifiable, so that higher risk leads to higher discount rates. Importantly, the controlled nature of our experiment allows us to demonstrate that the mechanism which gives rise to the variation in discount rates is distinct from common explanations based on systematic risk.

and Zafar, 2019; Das, Kuhnen, and Nagel, 2020; Choi and Robertson, 2020; Bordalo et al., 2020; Choi et al., 2021; Chinco, Hartzmark, and Sussman, 2021; Liu et al., 2021; De La O and Myers, 2021). These studies have spurred the development of new models, which aim to match not only data on realized returns, but data on return expectations from surveys (Barberis et al., 2015; Adam, Marcet, and Beutel, 2017). In our data, we observe a positive correlation between discount rates and expected cash flows, which is in line with the positive correlation between expected cash flow growth and expected returns documented in Giglio et al. (2021a). This suggests that new models may seek to additionally impose this covariation between expectations of fundamentals and returns (Jin and Sui, 2021)[4].

We also contribute to the literature on learning and asset pricing. In our design, subjects are told that there is a stock which follows a two-state Markovian process. In the good state, the cash flow distribution has a lower volatility and higher mean, compared to the cash flow distribution in the bad state. Importantly, the state can switch in each period with 20% probability. This switching process ensures that there is persistent time series variation in expected cash flows (see, for example, Veronesi, 2000; Lettau, Ludvigson, and Wachter, 2008; Collin-Dufresne, Johannes, and Lochstoer, 2016; Ghaderi, Kilic, and Seo, 2021). While the distribution of possible cash flows in our experiment does not have an extremely long left tail, we can still qualitatively examine how perception of downside risk affects discount rates. Moreover, because we provide subjects with the distribution of cash flows (conditional on each state), subjects are able to learn about downside risk even in the absence of a downside event.

[4] For related models that assume misperception of fundamentals, see Barberis et al. (1998), Fuster et al. (2012), Choi and Mertens (2013), Alti and Tetlock (2014); Hirshleifer et al. (2015), and Bordalo et al. (2018, 2019). For a more comprehensive review of these models and the associated empirical evidence, see Barberis (2018).

The rest of this paper is organized as follows. Section 1 presents our experimental design. We present our experimental results in Section 2. We discuss our results and conclude in Section 3.

1. Experimental Design

A. Experimental Setup

There is a stock that delivers a dividend, d_t, in each of 30 periods. There are five possible dividends: {\$60, \$85, \$115, \$135, \$150}, and the distribution of dividends is governed by a two-state Markov chain. We build a five point distribution of cash flows to mimic the five point distribution of returns that Giglio et al. (2021a) use to elicit beliefs from their survey respondents[5]. We denote the state in period *t* by s_t, which can take on one of two values, either *good* or *bad*. The dividend distribution generated in the good state has a higher mean and lower volatility, compared with the dividend distribution in the bad state. Specifically, in the bad state, the distribution of dividends is given by:

$$\Pr(d_t|s_t = bad) \equiv (\$60, 0.15; \$85, 0.30; \$115, 0.40; \$135, 0.10; \$150, 0.05). \tag{1}$$

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

$$\Pr(d_t|s_t = good) \equiv (\$60, 0.05; \$85, 0.10; \$115, 0.40; \$135, 0.30; \$150, 0.15). \tag{2}$$

[5] Giglio et al. (2021a) elicit a distribution over five different ranges of returns, whereas we will elicit a distribution over five different values of cash flows.

Note that the bad state is also associated with a higher probability of delivering the lowest dividend of \$60 (in addition to having higher volatility). We initialize the state in period 1 to be either good or bad with equal probability: $\Pr(s_1 = good) = 0.5$. There is also persistence in the states: 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 so that there is substantial time series variation in the expected dividend. Furthermore, the two-state switching process guarantees that the variation does not decline over time (as would be the case, in say, a model where there is probability 0 of switching from one state to the other). To ease comparability of behavior across subjects, we use the same realized sequence of thirty dividends for all subjects.

In eight randomly chosen periods, we elicit a subject's full distribution of expectations about next period's dividend. 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%. This elicitation enables us to test how the subjective distribution of dividend expectations differs from the objective distribution. We also ask subjects to report their willingness to pay for the right to receive next period's dividend, d_{t+1} (one can think of this as a willingness to pay for a one-period dividend strip). We use this willingness to pay to back out the implied discount rate, as described below.

Importantly, we incentivize the expectations question and the willingness to pay 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 quantity as $p_t = \Pr(s_t = bad \mid d_t, d_{t-1}, \dots d_1)$. Conditional on p_t, it is straightforward to compute the distribution of dividends.

Because the stochastic process is Markovian, we can rewrite the expression for p_t as a function of the current period's realized dividend and the prior belief:

$$p_t(p_{t-1}, d_t) =$$

$$\frac{\Pr(d_t|s_t = bad)\ \Pr(s_t = bad \mid p_{t-1})}{\Pr(d_t|s_t = bad)\ \Pr(s_t = bad|p_{t-1}) + \Pr(d_t|s_t = good)\ \Pr(s_t = good|p_{t-1})}$$

$$= \frac{\Pr(d_t|s_t = bad)\ (0.8p_{t-1} + 0.2(1 - p_{t-1}))}{\Pr(d_t|s_t = bad)\ (0.8p_{t-1} + 0.2(1 - p_{t-1})) + \Pr(d_t|s_t = good)\ (0.2p_{t-1} + 0.8(1 - p_{t-1}))} \tag{3}$$

where the expressions $\Pr(d_t|s_t = bad)$ and $\Pr(d_t|s_t = good)$ are defined in equations (1) and (2), respectively (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: $E(d_t|p_t) = p_t E[d_t|s_t = bad] + (1 - p_t)E[d_t|s_t = good]$. Similarly, the probability of each dividend outcome is a weighted average of the probability of that outcome in each of the two states. For example, for a \$60 dividend, $Pr(d_t = \$60|p_t) = p_t Pr(d_t = \$60|s_t = bad) + (1 - p_t)Pr(d_t = \$60|s_t = good)$. This establishes the Bayesian benchmark, which we use to incentivize subjects when they report their dividend expectations. In particular, we pay subjects based on their accuracy relative to the Bayesian benchmark (below we describe the payment scheme in more detail.)

It is more challenging to incentivize truthful reporting of the discount rate, which is a subjective quantity that depends on the subject's risk tolerance, and there is no natural Bayesian benchmark. Our solution is to ask subjects to report their willingness to pay for the right to receive next period's dividend, d_{t+1}. We then use this willingness to pay to infer the discount rate as follows:

$$DiscountRate_t = \frac{\mathrm{E}_{subjective}[d_{t+1}|p_t]}{WTP_t} \tag{4}$$

In the right hand side of (4), the numerator is the subjective expected next-period dividend, which we compute from the elicited distribution of dividends.

We incentivize subjects as follows. We randomly pick one of the eight periods in which we elicit the distribution of beliefs and the willingness to pay, and then pay subjects based on either the distribution question or the willingness to pay 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 are off, we subtract 3 cents.

If instead the willingness to pay 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 willingness to pay. 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 willingness to pay 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 willingness to pay, the subject purchases the one period asset at the randomly drawn price. If the number is larger than the stated willingness to pay, 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 USD. Thus, subjects can receive a bonus of up to $3 in this question.

While it may be difficult for subjects to implement the Bayesian updating rule in (3), our main tests do not rely critically on subjects' ability to accurately compute p_t. Rather, because we are interested in testing whether discount rates increase with risk, it is important to test this hypothesis using a measure of the subject's *perception* of the asset's risk. Thus, even if dividend expectations do not coincide with rational expectations, we can still test whether discount rates and perceived risk assessments are positively correlated.

As a rational benchmark, we can compute the relationship between discount rates and risk, assuming a subject is Bayesian and risk averse. We begin by defining risk as the conditional volatility of cash flow. At any point in time, the probability of being in the bad state, p_t, induces a mixture of the two cash flow distributions generated in the good and bad state. Our measure of objective risk is defined as the volatility of the time-varying mixture of cash flow distributions. In order to impose risk aversion, we assume the subject has a utility function given by u(x) = log(x). (In the Appendix, we show our benchmark predictions are robust to alternative utility functions that impose risk aversion).

Figure 1A plots the relationship between discount rates and risk, across the eight periods in which we elicit beliefs and WTP for the dividend strip in our experiment. There is a positive relationship: as volatility increases, a risk averse agent's willingness to pay drops, and she discounts cash flows at a higher rate – controlling for shifts in rationally forecasted cash flows. When deriving this relationship, we assume that the agent views the asset's cash flow risk as being non-diversifiable, and thus she demands compensation for this risk in the form of high expected returns. Of course, in a rational framework, the presence of only diversifiable risk would induce a flat relationship between discount rates and risk. As we will see, the data from our experiment closely mimics data from surveys about the aggregate stock market. In this sense, we believe that the mechanism generating the data in our experiment – despite the absence of systematic risk – has implications for the behavior of investors in the field. In Figure 1B, we show that the same positive relationship holds when using an alternative measure of risk based on the probability of the lowest dividend. Because we have time series variation of discount rates and willingness to pay within subject, we can test for this relationship at the subject level.

Figure 1A: Discount rate vs. risk for a Bayesian risk averse agent, using conditional volatility as the measure of risk.

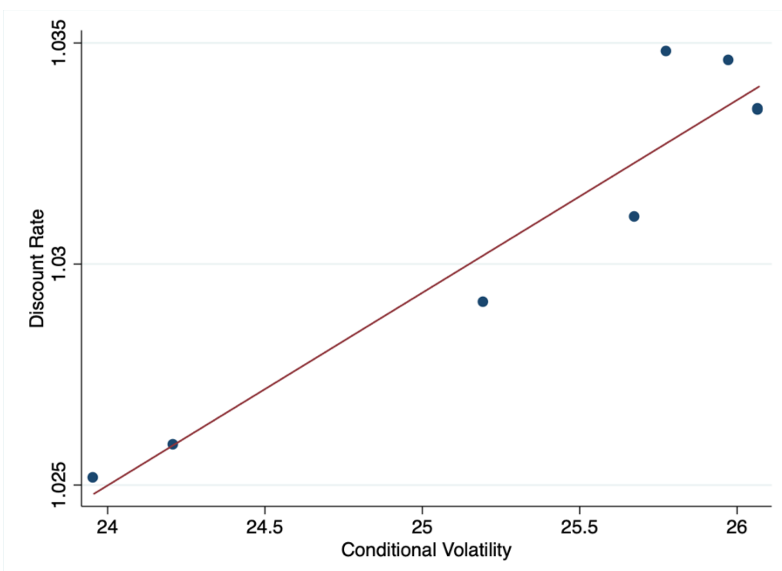

Figure 1B: Discount rate vs. risk for a Bayesian risk averse agent, using the probability of the lowest dividend as the measure of risk.

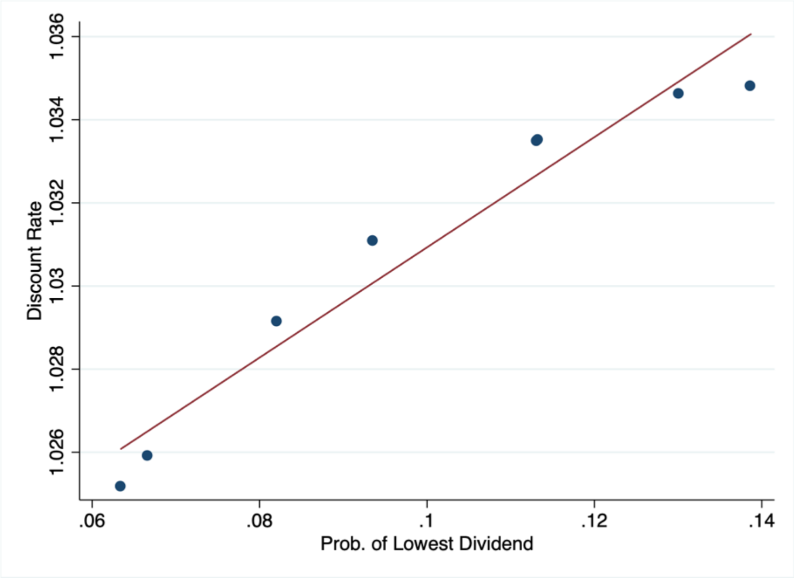

B. *Experimental Procedures*

We recruit N=300 subjects from the online data collection platform, Prolific. The sample size, main hypotheses, and exclusion criteria are all pre-registered on Aspredicted.org (see https://aspredicted.org/6Z4_RLQ for the pre-registration document). Subjects receive $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 earning is $4.39, including the $2 participation fee.

2. *Experimental Results*

A. *Sample and Summary Statistics*

In Figure 2A we plot the time series of expected dividends and WTP. At each elicitation period, the subjective expected dividend, E*[D], is the mean of the subject's reported cash flow distribution. The figure displays the average across all subjects at each of the eight elicitation periods. Several things stand out in this graph. First, subjects' WTP is consistently below their expected dividend, which implies that subjects are risk averse on average. Second, WTP follows a similar path as E*[D], suggesting that subjects adjust their WTP according to changes in their beliefs.

In Figure 2B, we additionally overlay our measure of risk, namely the volatility implied by subjects' reported cash flow distribution.[6] The figure shows that both E*[D] and WTP are strongly negatively correlated with risk in the time series. In some tests, we use the perceived probability of the lowest dividend as our measure of risk, in analogy to disaster risk models, such as Wachter (2013). Figure 2C shows that the two measures of risk are strongly positively correlated in the time series. Indeed, all our results are robust to using either measure of risk. In Table 1, we provide additional summary statistics for our main variables.

[6] To ease interpretation, we add a second y-axis for the perceived conditional volatility. Due to the two different scales of the y-axis, we omit the confidence intervals.

Figure 2A: This figure shows the average expected dividend E*[D] and average WTP across all subjects at each of the eight elicitation periods. The vertical bars denote 95% confidence intervals, clustered by subject.

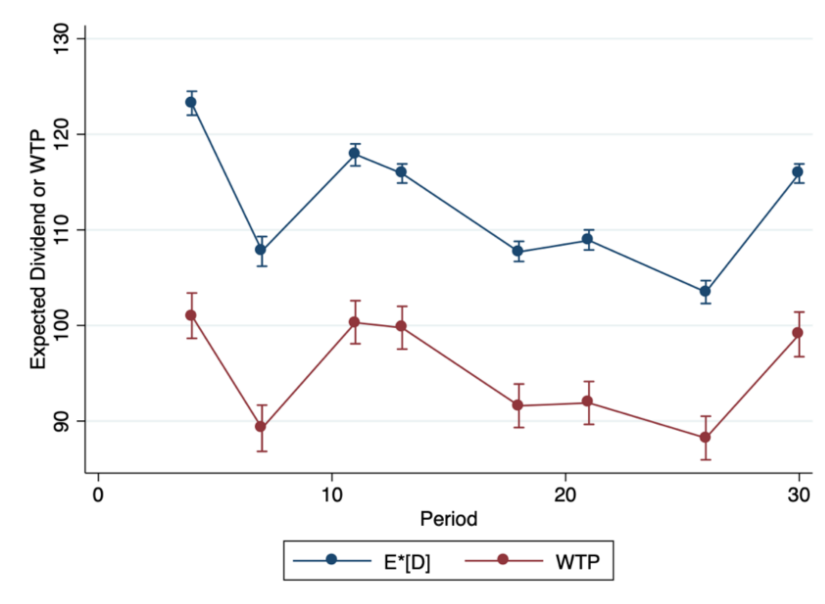

Figure 2B: This figure shows the average expected dividend E*[D], average WTP, and average perceived conditional volatility across all subjects at each of the eight elicitation periods.

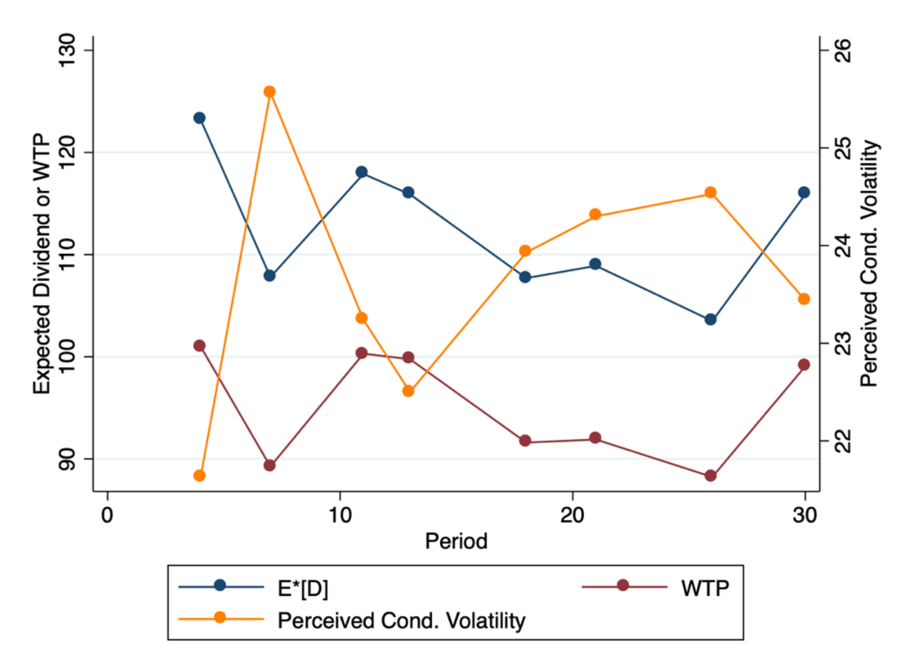

Figure 2C: This figure shows the average perceived conditional volatility and average perceived probability of the lowest dividend across all subjects at each of the eight elicitation periods.

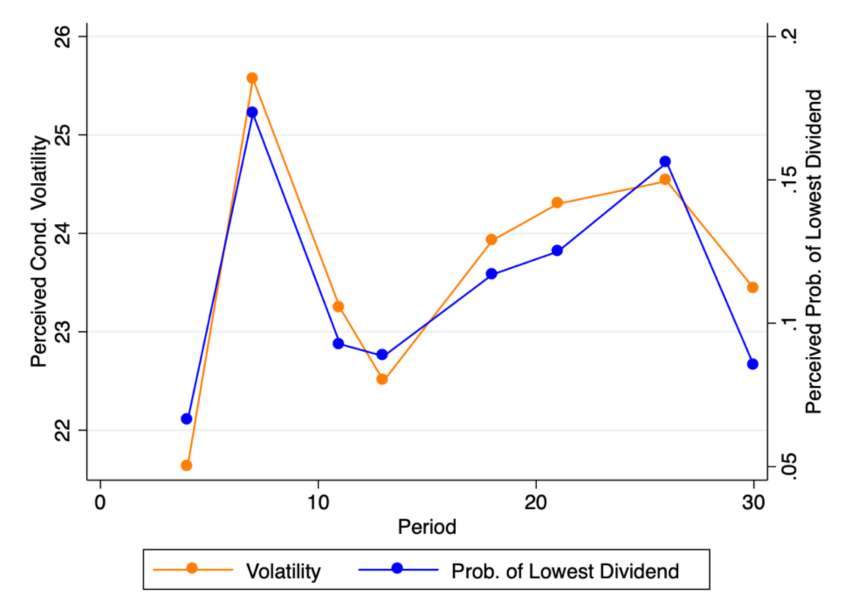

Table 1: Summary Statistics

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. E*[D] is the subjective expected dividend, defined as the mean of a subject's reported distribution. E[D] is the Bayesian expected dividend, defined as the mean of the Bayesian distribution. WTP is the subject's reported willingness to pay for next period's dividend. Discount rate is the ratio of E*[D] and WTP. Expected realized return is the ratio of E[D] and WTP. Perceived conditional volatility is the volatility of a subject's reported cash flow distribution. Perceived probability of lowest dividend is the subjective probability that a subject attaches to the lowest dividend in the reported distribution.

	Mean	**p25**	**p50**	**p75**	**Std. Dev.**	**Min**	**Max**	**N**
Expected Dividend E*[D]	112.61	105.50	113.00	120.50	12.04	65.00	150.00	2,400
WTP	95.15	80.00	95.70	110.00	20.88	60.00	150.00	2,400
Discount Rate	1.23	1.04	1.19	1.38	0.27	0.52	2.23	2,400
E*[D] / E[D]	1.01	0.96	1.02	1.08	0.10	0.59	1.32	2,400
Expected Realized Return	1.22	1.02	1.17	1.38	0.27	0.69	1.97	2,400
Perceived Conditional Volatility	23.64	21.36	24.81	27.22	6.05	0.00	39.69	2,400
Perceived Prob. of Lowest Dividend	0.11	0.05	0.10	0.15	0.11	0.00	0.90	2,400

B. Discount Rates and Risk

In Figure 3 we plot the empirical relationship between discount rates and risk. We define risk as perceived conditional volatility, which we compute as the volatility of the cash flow distribution reported by the subject. The discount rate is defined in (4). Figure 3 displays a clear negative relationship between perceived risk and discount rates. This is starkly at odds with the rational benchmark plotted in Figure 1, which predicts a strong *positive* relationship between discount rates and risk. To be clear: when subjects perceive risk to be higher, we find that subjects apply less of a discount to future cash flows. We formally confirm this relationship in column 1 of Table 2, which present results from a mixed effects regression with a random slope and a random intercept. This specification allows for heterogeneity across subjects in average discount rates, and also heterogeneity with respect to the sensitivity between discount rates and perceived risk.

While the data strongly reject the prediction of the rational benchmark, the results are consistent with the evidence from Giglio et al. (2021a), who find that expected returns are negatively correlated with perceived risk. Because we measure both valuations and expectations, we can use these two components of the discount rate to analyze the mechanism that generates the negative relationship in Figure 3.

Figure 3: Binned scatterplot of inferred discount rates and perceived risk, controlling for subject fixed effects.

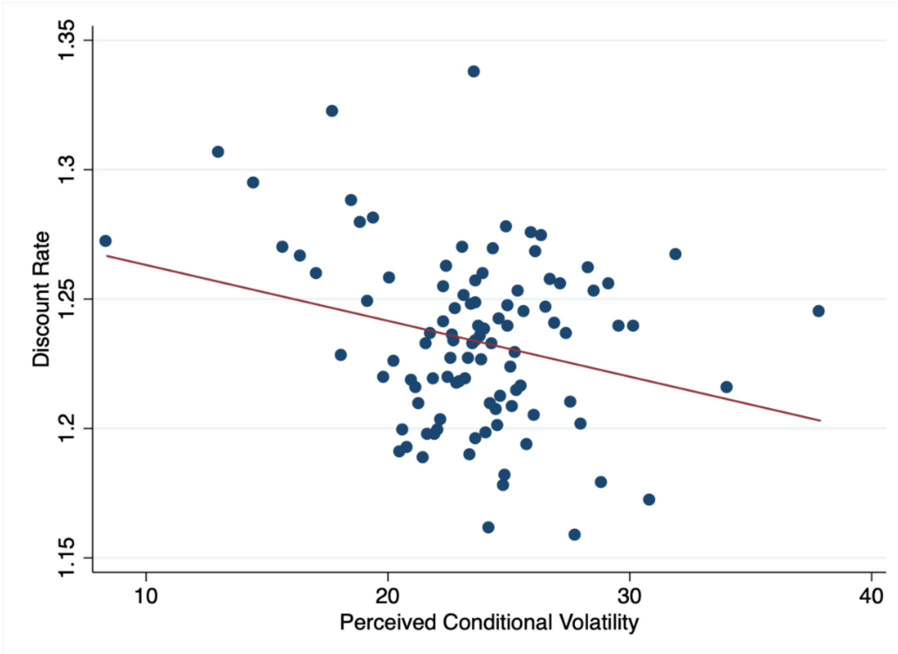

One candidate mechanism operates through a subject's willingness to pay. In particular, if subjects have rational expectations about dividends, but are willing to pay *more* for the asset when risk is higher, then this could give rise to a negative relationship between risk and discount rates. To test this hypothesis, in Figure 4 we plot the relationship between willingness to pay and perceived risk. The plot clearly shows that subjects are willing to pay less for the asset as risk increases. We formally confirm this result in a mixed effects regression in column 2 of Table 2.

Taking stock, so far we have found that both discount rates and willingness to pay are declining in perceived risk. The only way to reconcile these facts with the present value relationship in (4), is to allow cash flow expectations to depart from rational (Bayesian) expectations. Because we control the information that subjects have about the cash flow process, we can precisely measure how close (or far) subjective expectations are from the rational benchmark. To do so, we construct a variable which measures the degree to which subjects overestimate cash flows. This variable is given by the ratio, E*[D]/E[D], where E*[D] is the mean of the subjective cash flow distribution. Thus, for a Bayesian, E*[D]/E[D] should equal 1 in all periods, and should not vary with risk. Figure 5 shows that instead, E*[D]/E[D] declines systematically with risk. This clearly rejects rational expectations, and we confirm this negative relationship in column 3 of Table 2.

The negative relationship between E*[D]/E[D] and risk provides further evidence on the nature of the irrational expectations. When E*[D]/E[D] > 1, subjects overestimate cash flows and when E*[D]/E[D] < 1, subjects underestimate cash flows. Therefore, the data indicate that when risk is low, subjects are too optimistic; when risk is high, subjects are too pessimistic. Cash flow expectations are therefore moving too much, relative to a given change in perceived risk.

Figure 4: Binned scatterplot of willingness to pay and perceived risk, controlling for subject fixed effects.

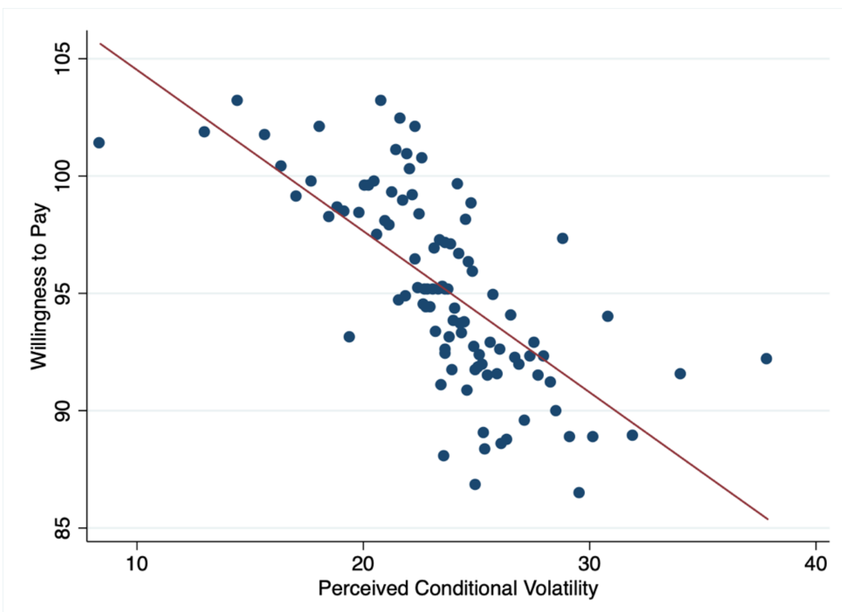

Figure 5: Binned scatterplot of E*[D]/E[D] and perceived risk, controlling for subject fixed effects. The horizontal line at E*[D]/E[D]=1 provides the rational benchmark which predicts no (variation in) overestimation of cash flows.

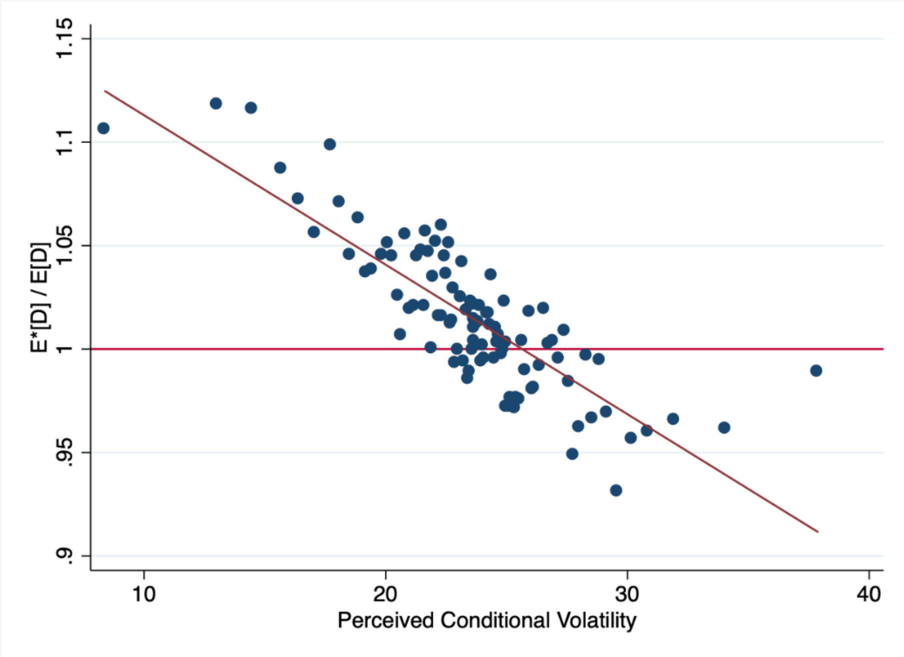

This excessive volatility of beliefs can naturally explain the negative relationship between discount rates and risk in Figure 3. Since we know that WTP is declining in risk, then subjective dividend expectations must be falling faster than WTP, in order for the ratio of E*[D]/WTP to decline in risk. The excessive volatility of beliefs generates the excessive sensitivity of E*[D] to risk. Our data therefore indicate that when risk increases, subjects discount cash flows at a lower rate because of wrong cash flow beliefs. Interestingly, the excessive volatility of beliefs that we observe does not necessarily imply that on average, beliefs are wrong. The fact that we elicit a time series of expectations within subject is crucial to identifying the source of the irrational expectations, which we examine further in the next section.

Figure 6 shows the relationship between discount rates and expectations of cash flows. When subjects expect high cash flows, they apply high discount rates; conversely, when subjects expect low cash flows, they apply low discount rates. This finding is in stark contrast to the Bayesian benchmark, which predicts a negative relationship. However, the data are consistent with Giglio et al. (2021a), who find that expected returns and expected GDP growth (a proxy for cash flow growth) are positively correlated.

Figure 6: Binned scatterplot of inferred discount rates and E*[D], controlling for subject fixed effects.

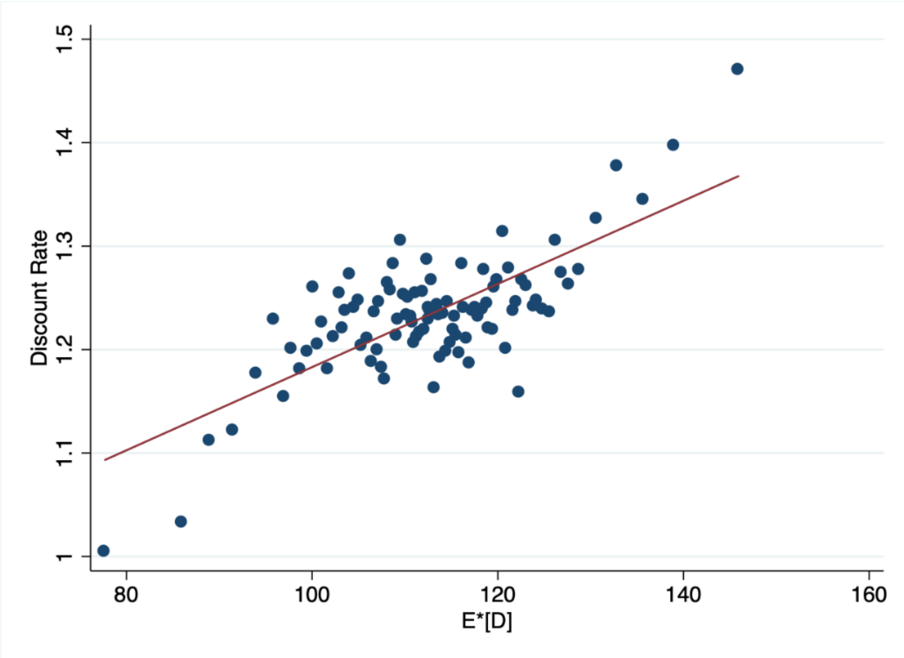

C. Expected Realized Returns and Risk

We can also test whether the expected *realized* returns increase in risk – as they do in the field. To do so, we replace the subjective expectation in the numerator in (4) with rational expectations. Figure 7 presents a clear picture: there is a strong positive relationship between perceived risk and average realized returns (see column 4 of Table 2 for regression results). When subjects set the price of the asset low (in times of high risk), subsequent realized returns are high. Thus, we reproduce two key facts from the field: subjective expected returns vary negatively with risk, while expected realized returns vary positively with risk (Greenwood and Shleifer, 2014).

One concern, however, with interpreting the survey evidence from Greenwood and Shleifer (2014) through the lens of representative agent asset pricing models, is that survey respondents may not be the marginal investor. Our design sidesteps this concern because we elicit both expectations *and* valuations from each respondent. Thus, we are able to examine the endogenous relationship between realized returns and expectations at the individual level, without having to appeal to general equilibrium pricing.

Figure 7: Binned scatterplot of expected realized returns and perceived risk, controlling for subject fixed effects.

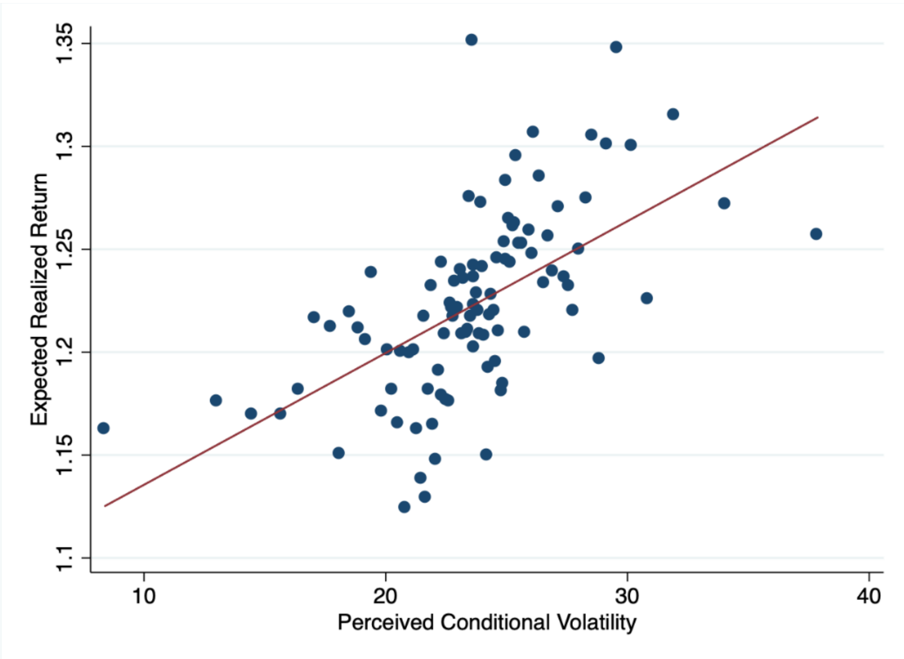

Table 2: Univariate regression results with two alternative definitions of perceived risk.

This table presents results from a mixed effects regression with a random slope and a random intercept in each univariate specification. The dependent variable in column 1 is a subject's inferred discount rate, in column 2 it is a subject's willingness to pay, in column 3 it is the ratio of a subject's expected dividend to the Bayesian expected dividend, and in column 4 it is a subject's expected realized return. The independent variable in Panel A is a subject's perceived conditional volatility, and in Panel B it is a subject's perceived probability of the lowest dividend. Standard errors are clustered by subject and displayed in parentheses below the coefficient. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A	(1)	(2)	(3)	(4)
Dependent variable:	Discount rate	WTP	E*[D] / E[D]	Expected realized return
Perceived conditional volatility	-0.002**	-0.662***	-0.006***	0.006***
	(0.001)	(0.079)	(0.000)	(0.001)
Observations	2,400	2,400	2,400	2,400

Panel B	(1)	(2)	(3)	(4)
Dependent variable:	Discount rate	WTP	E*[D] / E[D]	Expected realized return
Perceived prob. of lowest dividend	-0.253***	-67.529***	-0.665***	0.671***
	(0.056)	(5.039)	(0.016)	(0.062)
Observations	2,400	2,400	2,400	2,400

D. Overestimation and Underestimation of Cash Flows

Because we elicit the full distribution of cash flow beliefs, we can further analyze whether expectational errors are concentrated on specific cash flows and how these expectational errors vary with perceived risk[7]. In Figure 8A, we plot the subjective distribution of beliefs against the objective distribution. Because the objective distribution is time-varying, we average it across all eight elicitation periods. We see small deviations from the rational benchmark, mainly in the form of an underestimation of the $115 dividend and an overestimation of the smaller probability $135 and $150 dividends. On net, this leads to a slight overestimation of mean cash flows (subjective: $112.61 vs. objective: $110.97, p<0.001). This is consistent with the evidence from Figure 5, which demonstrates that for a bulk of our data, E*[D]/E[D]>1, implying overestimation of cash flows.

At the same time, we know from Figure 5 that deviations from rational expectations are systematically related to perceived risk. Specifically, subjects overestimate expected cash flows to a greater extent as perceived risk declines. Here we investigate this pattern more closely by examining the full distribution of beliefs. In Figure 8B, we plot the difference between subjective and objective beliefs, and we cut the data based on whether perceived volatility falls above or below the sample median.

We see that when subjects perceive risk to be low, they underestimate the probability of the two lowest cash flows. This generates an overestimation of the mean cash flow, relative to the Bayesian benchmark (subjective: $116.44 vs. objective: $111.70, p<0.001). Conversely, when subjects perceive high risk, they *overestimate* the probability of the lowest cash flow,

[7] De La O and Myers (2021) provide some evidence that analyst forecast errors about earnings are predictable in some subsamples, which suggests overreaction. We complement their analysis by inspecting the mechanism in a fully controlled environment, allowing us to study the nature of expectation errors conditionally.

pushing the subjective expected cash flow below the Bayesian benchmark (subjective: \$108.77 vs. objective: \$110.26, $p<0.001$). It is worth emphasizing the quantitative nature of subjects' distorted beliefs about the lowest cash flow. When perceived risk is low, the objective Bayesian probability of the lowest cash flow is 9.6%; subjects state the subjective probability is 6.4%, leading to an approximate distortion of 33%. When perceived risk is high, the objective Bayesian probability of the lowest cash flow increases to 10.4%; in this case, subjects state the subjective probability is 16.1%, leading to an approximate distortion of over 50%. Taken together, the evidence suggests that subjective expected cash flows move too much for a given shift in perceived risk; furthermore, this excess volatility of beliefs stems in large part from wrong beliefs about the lowest possible cash flow.

One caveat to the above analysis is that we are analyzing how belief distortions vary with perceived risk, rather than objective risk. Thus, the degree to which expected cash flows are "excessively" volatile is computed against the benchmark of how much expected cash flows should move for a given change in perceived risk. At the same time, the above analysis is motivated by the fact that measured beliefs are not rational, and thus perceived risk will not coincide with objective risk. Thus, the "excessive volatility" that we observe in cash flow expectations is not identical to the excessive volatility in the sense of Shiller (1981).

Figure 8A: Subjective and objective distribution of cash flow expectations. 95% confidence intervals are included and standard errors are clustered by subject.

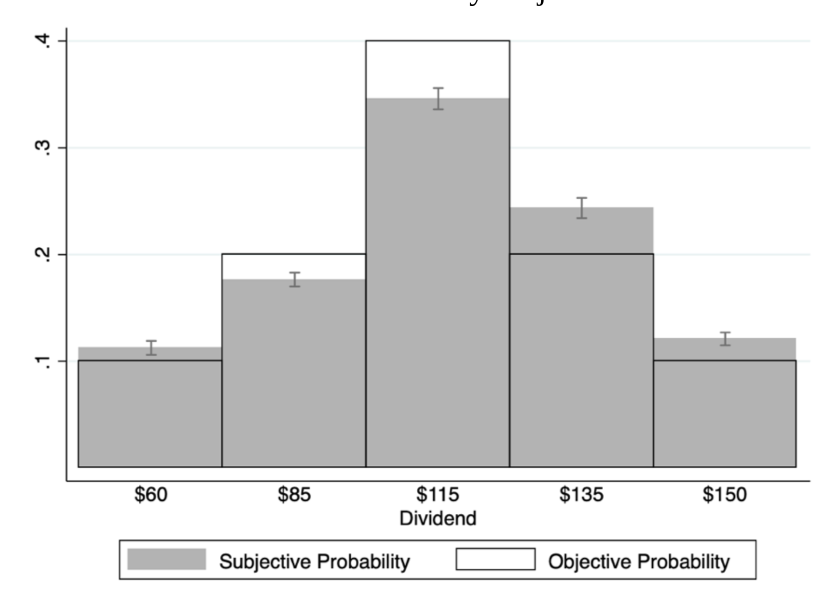

Figure 8B: Deviations from Bayesian cash flow beliefs. Data are cut by whether the belief is associated with perceived volatility that is above the median (dark grey bars) or below the median (light grey bars).

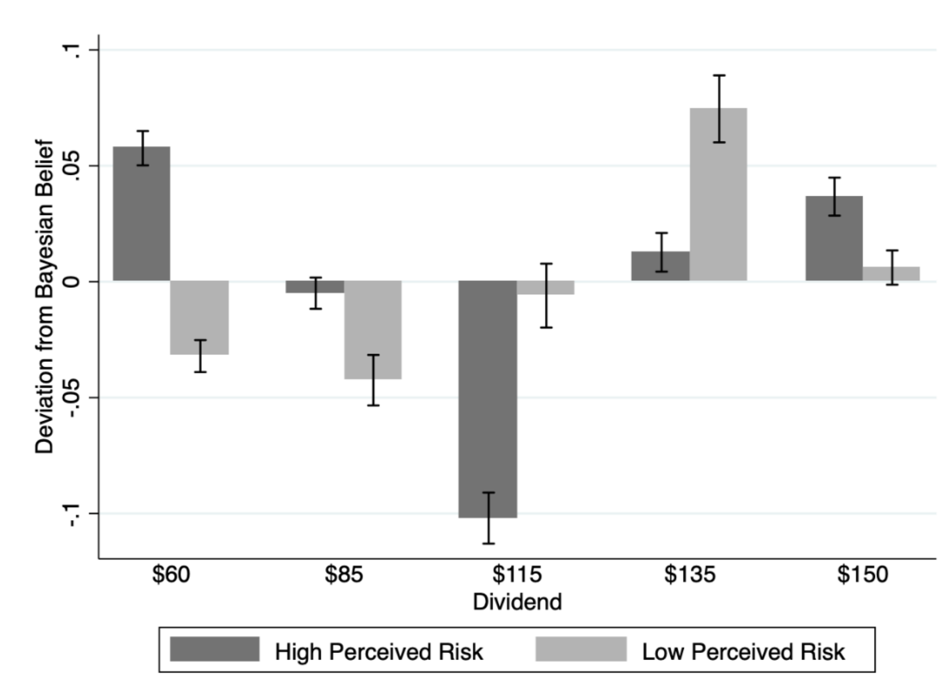

If we look back to Figure 5, it is important to emphasize that those results indicate that the systematic relationship between wrong cash flow beliefs and perceived risk obtains within subject (since we include subject fixed effects.) Yet it is possible that when partitioning the data based on high and low levels of perceived risk, as in Figure 8B, this could mainly partition the data *across* subjects. In other words, those observations where subjects perceive risk to fall below the sample median may be concentrated among a group of persistently optimistic subjects. The observations where subjects perceive risk to be high (above sample median) would then be concentrated among a group of pessimistic subjects. In the next section, we investigate the degree to which some subjects have persistently low or high expectations of cash flows.

E. Heterogeneity Across Subjects

For each subject, we compute the subjective expectation of cash flows, averaged across the eight elicitation periods. If all subjects hold rational expectations about cash flows, then the average subjective cash flow expectation would equal $110.97. Figure 9A 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, on average, pessimists who expect cash flows below the rational expectation; there are also many subjects who are optimists and expect cash flows above the rational expectation.

Figure 9B conveys a similar message about perceived volatility, there are some subjects who on average perceive the cash flow distribution to be riskier than other subjects. This substantial heterogeneity is consistent with the large investor fixed effects found in survey data (Giglio et al. 2021a). Figure 9C shows that it is the pessimists – those who on average expect low cash flows – who are more likely to perceive high risk. This joint distribution across subjects complements the within-subject effect that we display in Figure 5[8].

Figure 9A: Kernel density of subjective expectations across subjects. For each subject, we compute the average subjective expectation of the next-period dividend across the eight elicitation periods. The plot shows the density across subjects. The red vertical line shows the average cash flow expectation of a Bayesian investor.

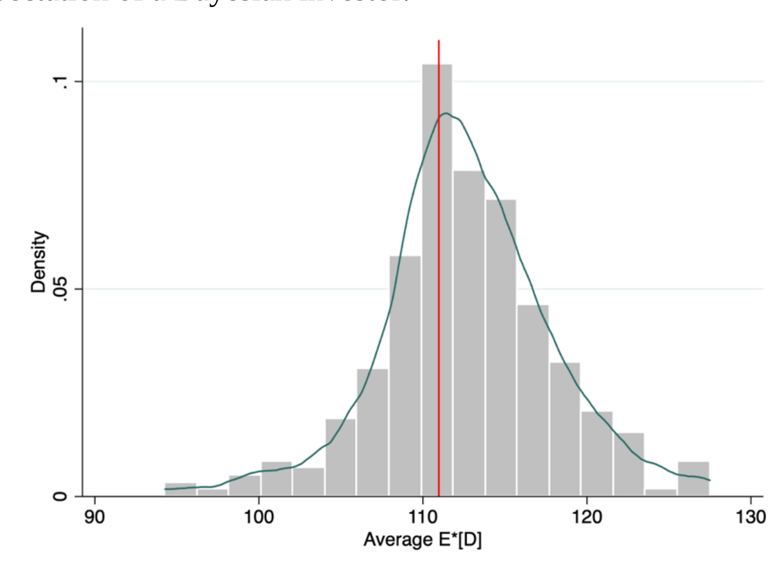

[8] We note that Figure 5 and Figure 9C are not perfectly comparable. In Figure 5, we display how E*[D]/E[D] declines in perceived risk; Figure 9C displays how E*[D] declines in perceived risk.

Figure 9B: Kernel density of perceived risk across subjects. For each subject, we compute the average perceived risk across the eight elicitation periods. The plot shows the density across subjects. The red vertical line shows the average conditional volatility of a Bayesian investor.

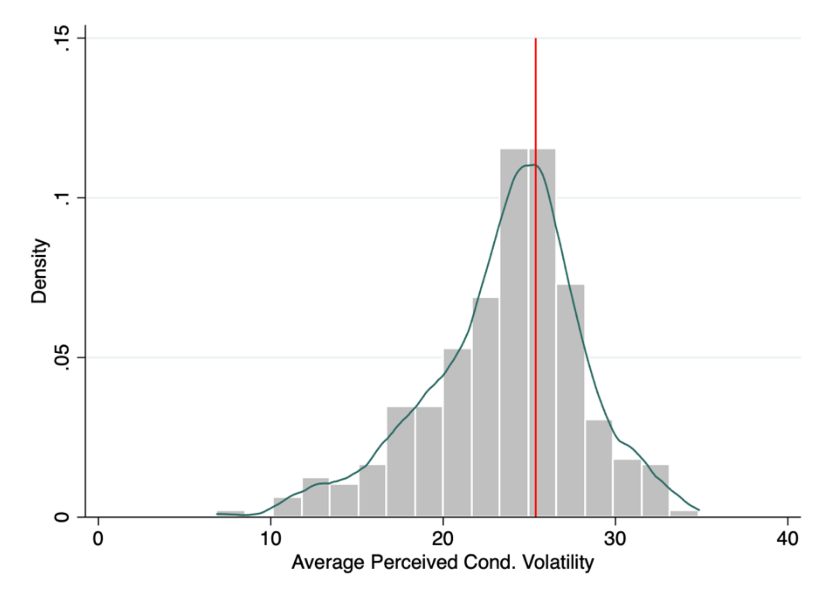

Figure 9C: Scatterplot of perceived risk vs. subjective expected cash flow. Each point represents a single subject, and is computed by averaging perceived conditional volatility and E*[D] across the eight elicitation periods.

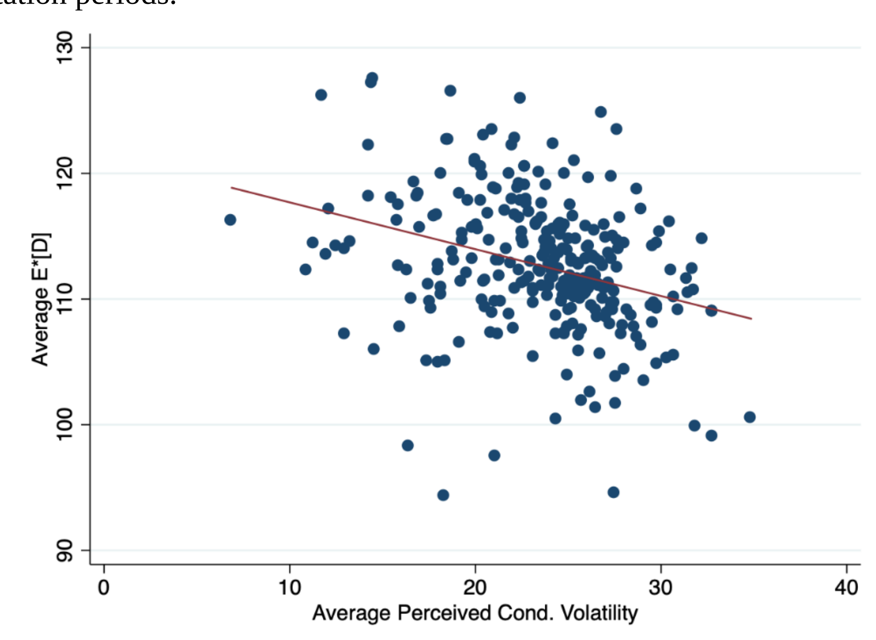

3. Discussion

We have conducted an experiment to analyze several facts from the literature on survey expectations. Just as in the surveys, we elicit measures of perceived risk and expected returns. The novel experimental design that we employ, however, moves beyond a traditional survey design by incorporating several advantages, including precise control over subjects' information sets and the incentive compatible elicitation of discount rates.

Our experimental results stand in stark contrast to the predictions of a model in which a rational risk averse agent forms expectations using Bayes' rule. At the same time, several patterns in our data resemble features of expectations and portfolio choice from a sample of Vanguard investors (Giglio et al. 2021a; 2021b). Throughout the paper, the main pattern we have emphasized is the negative correlation between discount rates and perceived risk.

In addition, we also find substantial heterogeneity in expectations across subjects, similar to the substantial investor fixed effects found in Giglio et al. (2021a). Furthermore, we document that expected returns and expected cash flows are positively correlated. This is reminiscent of the covariation between expected returns and expected cash flow *growth* (proxied by GDP growth) found in Giglio et al. (2021a). Our ability to experimentally reproduce these facts using a substantially different methodology, should provide confidence that data from previous surveys are largely immune to critiques about lack of incentives or confusion about expected returns.

The empirical correlation between expected returns and expected cash flows is related to the model of Jin and Sui (2021), which predicts covariation between the two expectations. In that model, by assumption investors extrapolate past stock market returns: expected returns are positively related to sentiment, which itself is an increasing function of past returns. Importantly, the model endogenously generates expectations about cash flow growth that are *more* sensitive to sentiment, compared with expectations of returns; this is due to the mean reversion in the price-dividend ratio. In this sense, the experimental data we produce in a one-period setting are broadly consistent with the prediction of the Jin and Sui (2021) model -- if one replaces sentiment with risk perception. In particular, we find that subjects' discount rate declines in risk, yet their cash flow expectation declines faster than discount rates, leading to low prices in times of high risk. Interestingly, some of the forces in the Jin and Sui (2021) model which induce covariation between cash flow expectations and return expectations are absent in our experimental setting. Given that we still observe a stronger reaction of cash flow expectations to the underlying state compared to subjective return expectations, our findings bolster the case for new models to maintain this relation.

A key message in our paper is that the negative relationship between discount rates and perceived risk is generated by non-Bayesian cash flow expectations. In particular, Figure 5 shows that subjects overestimate cash flows when perceived risk is low, but they switch to underestimating cash flows when perceived risk is high. It is therefore crucial to measure and control the perception of risk, otherwise we would wrongly conclude that cash flow expectations are on average correct but subject to idiosyncratic noise.

To see this point more clearly, we reproduce Figure 2A below by overlaying the time series of the Bayesian expectation of cash flows over subjective expectations. The picture suggests that on average, subjects track the objective expected dividend quite well. How can this picture be reconciled with the clear deviations from rational expectations in Figure 5? It turns out that at a given point in time, there is substantial heterogeneity across subjects in perceived risk, and importantly, subjective cash flow expectations are declining perceived risk. This analysis suggests caution when examining aggregate expectations data (Bordalo et el., 2020), and reinforces the importance of investor fixed effects emphasized by Giglio et al. (2021a).

Figure 10: This figure reproduces Figure 2A and shows the average expected dividend E*[D] and average WTP across all subjects at each of the eight elicitation periods. This figure also includes the Bayesian expected dividend E[D]. The vertical bars denote 95% confidence intervals, clustered by subject.

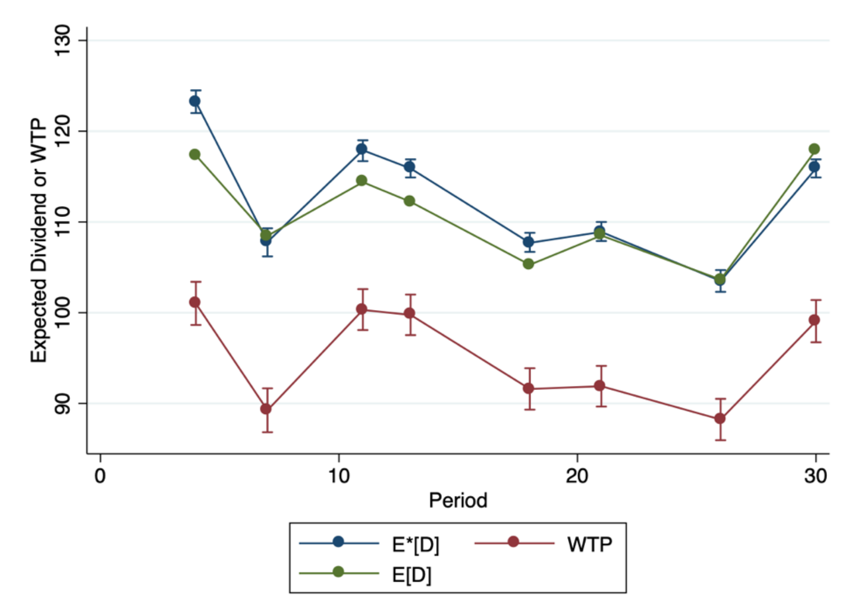

We have argued that the one period nature of the asset in our experiment is useful because it allows us to see how valuation relates to expectations in the simplest possible setting. Indeed, we find clear structure in subjective expectations, even when there is no need to form expectations over long horizons. Yet this simplicity also means that the data set we produce is not optimized to analyze other previously documented facts about subjective expectations.

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, Huang, and Jin, 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. 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.

References

Adam, K., A. Marcet, and J. Beutel, 2017, Stock price booms and expected capital gains. *American Economic Review* 107, 2352-2408.

Adam, K., D. Matveev, and S. Nagel, 2021, Do survey expectations of stock returns reflect risk adjustments?, *Journal of Monetary Economics* 117, 723-740.

Afrouzi, H., Kwon, S., Landier, A., Ma, Y. and D. Thesmar, 2021, Overreaction in expectations: Evidence and theory, *Working paper.*

Alti, A. and P. Tetlock, 2014, Biased beliefs, asset prices, and investment: a structural approach, *Journal of Finance* 69, 325–361.

Bansal, R. and A. Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

Barberis, N., 2018, Psychology-based models of asset prices and trading volume, *Handbook of Behavioral Economics: Applications and Foundations* 1, 79-175, North-Holland.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1-24.

Barberis, N., A. Shleifer, and R. Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–345.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer, 2019, Diagnostic expectations and stock returns, *Journal of Finance*, 74, 2839-2874.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer, 2020, Overreaction in macroeconomic expectations, *American Economic Review* 110, 2748-82.

Bordalo, P., N. Gennaioli, and A. Shleifer, 2018, Diagnostic expectations and credit cycles, *Journal of Finance* 73, 199-227.

Brunnermeier, M., E. Farhi, R. Koijen, A. Krishnamurthy, S. Ludvigson, H. Lustig, S. Nagel, and M. Piazzesi, 2021, Perspectives on the future of asset pricing, *Review of Financial Studies* 34, 2126–2160.

Campbell, J. and J. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205-251.

Campbell, J. and R. Shiller, 1988, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661-676.

Chinco, A., S. Hartzmark, and A. Sussman, 2021, A new test of risk factor relevance, *Journal of Finance*, forthcoming.

Choi, J., S. Bender, D. Dyson, and A. Robertson, 2021, Millionaires speak: What drives their personal investment decisions?, *Journal of Financial Economics*, conditionally accepted.

Choi, J. and T. Mertens, 2013, Extrapolative expectations and the equity premium, *Working paper*.

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

Cochrane, J., 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047-1108.

Collin-Dufresne, P., M. Johannes, and L. Lochstoer, 2016, Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review* 106, 664-98.

Da, Z., X. Huang, and L. Jin, 2021, Extrapolative beliefs in the cross-section: What can we learn from the crowds?, *Journal of Financial Economics* 140, 175-196.

Das, S., C. Kuhnen, and S. Nagel, 2020, Socioeconomic status and macroeconomic expectations, *Review of Financial Studies* 33, 395-432.

De La O, R., and S. Myers, 2021, Subjective cash flow and discount rate expectations, *Journal of*

Finance 76, 1339-1387.

Frydman, C., N. Barberis, C. Camerer, P. Bossaerts, and A. Rangel, 2014, Using neural data to test a theory of investor behavior: An application to realization utility, *Journal of Finance* 69, 907-946.

Fuster, A., B. Hebert, and D. Laibson, 2012, Natural expectations, macroeconomic dynamics, and asset pricing, NBER Macroeconomics Annual 26, 1–48.

Gabaix, X., 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645-700.

Ghaderi, M., M. Kilic, and S. Seo, 2021, Learning, slowly unfolding disasters, and asset prices. *Journal of Financial Economics,* forthcoming.

Giglio, S. and B. Kelly, 2018, Excess volatility: Beyond discount rates. *Quarterly Journal of Economics,* 133, 71-127.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus, 2021a, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481-1522.

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

Greenwood, R. and A. Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714-746.

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

Jin, L. and P. Sui, 2021, Asset pricing with return extrapolation, *Journal of Financial Economics,* forthcoming.

Kuchler, T. and B. Zafar, 2019, Personal experiences and expectations about aggregate outcomes, *Journal of Finance* 74, 2491-2542.

Lettau, M., S. Ludvigson, and J. Wachter, 2008, The declining equity premium: What role does macroeconomic risk play?, *Review of Financial Studies* 21, 1653-1687.

Liu, H., C. Peng, W. A. Xiong, and W. Xiong, 2021, Taming the bias zoo, *Journal of Financial Economics*, forthcoming.

Shiller, R., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421-436.

Veronesi, P., 2000, How does information quality affect stock returns?, *Journal of Finance* 55, 807-837.

Wachter, J., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987-1035.

Appendix

Figure 1 in the main text shows that for a risk averse agent with $u(x) = \log(x)$, the relationship between the discount rate and risk is positive. Here, we show that this relationship is robust to alternative utility functions that impose risk aversion. The following figures plot the linear slope coefficient that emerges in Figure 1 for different utility functions of the form $u(x) = x^{\alpha}$. The x-axis displays different values of α in [0.5, 1], and the y-axis displays the corresponding slope coefficient that emerges in Figure 1. The slope is positive for all values of α, showing that the relationship between discount rates and risk is firmly positive. As α increases, the agent becomes less risk averse, and the slope coefficient becomes smaller. For $\alpha = 1$, the agent is risk neutral, the relationship between the discount rate and risk is zero, and consequently the slope coefficient is zero. Mirroring Figure 1 from the main text, in Figure A.1A, we use the conditional volatility of the cash flow distribution as the measure of risk, and in Figure A.1B, we use the probability of the lowest dividend as the measure of risk.

Figure A.1A

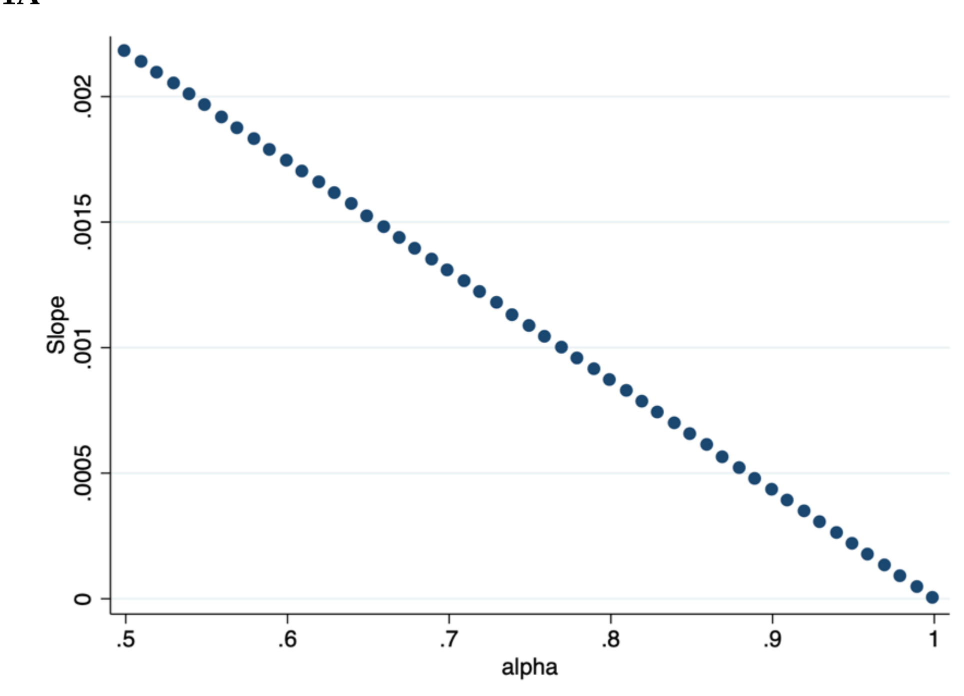

Figure A.1B

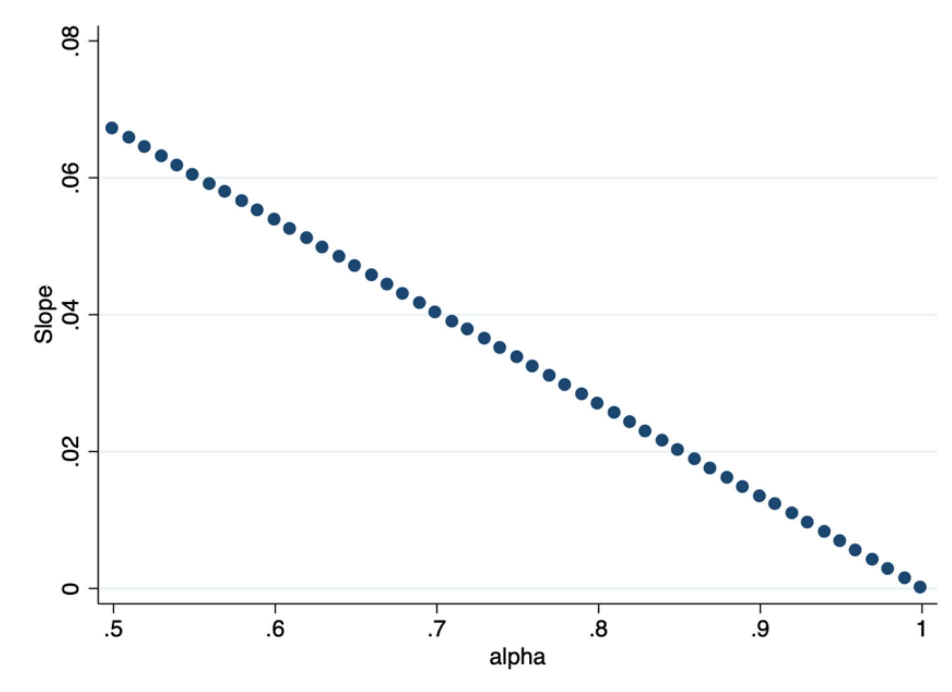

Screenshots of the Experiment

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. Below are screenshots showing how the dividend distributions were displayed to subjects:

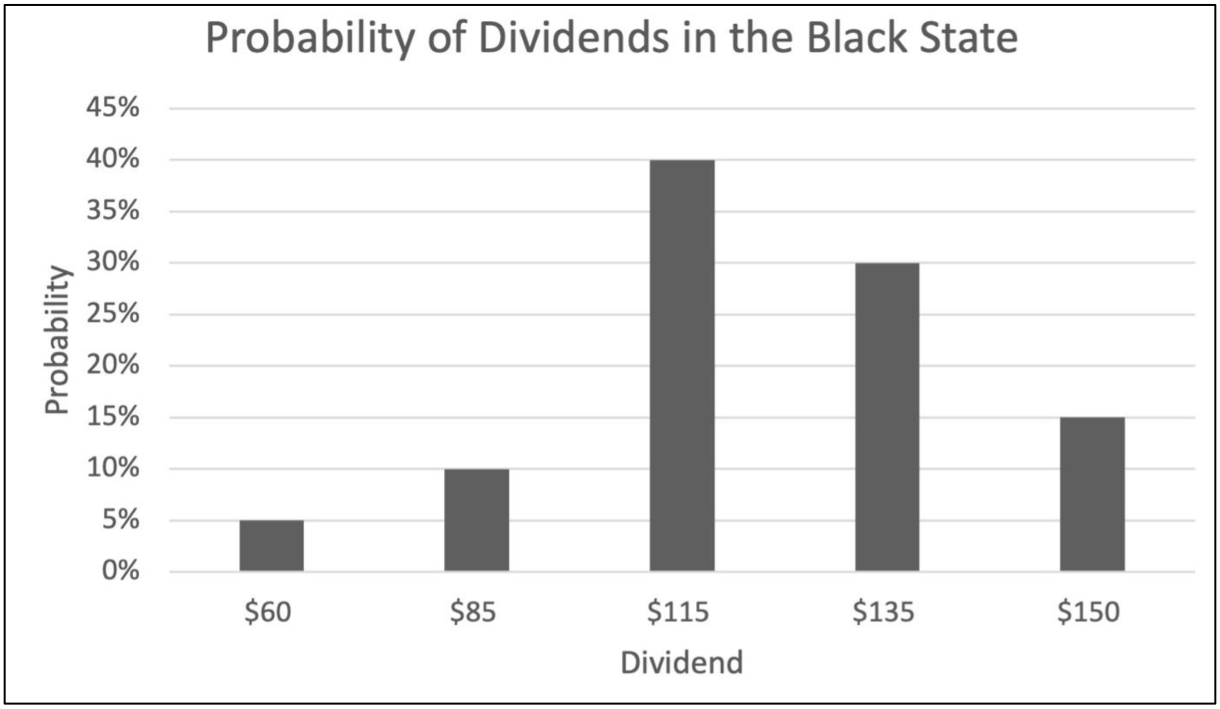

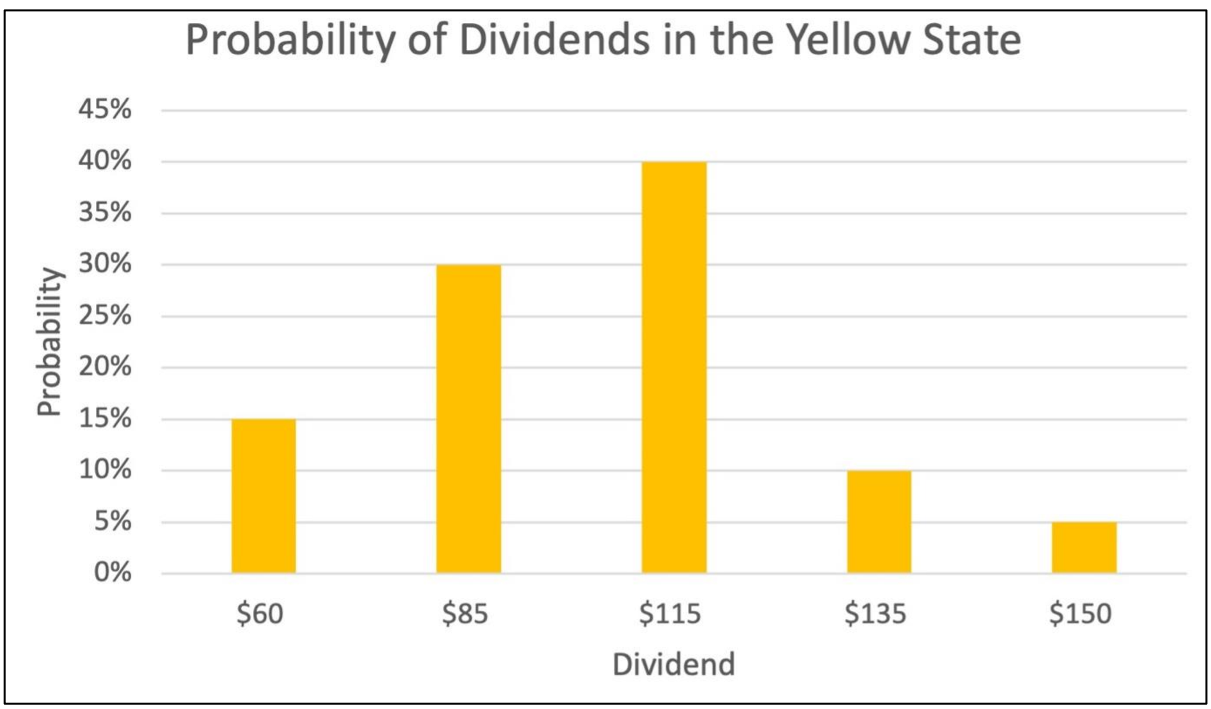

Below is a screenshot showing how dividend realizations were displayed to subjects in each period:

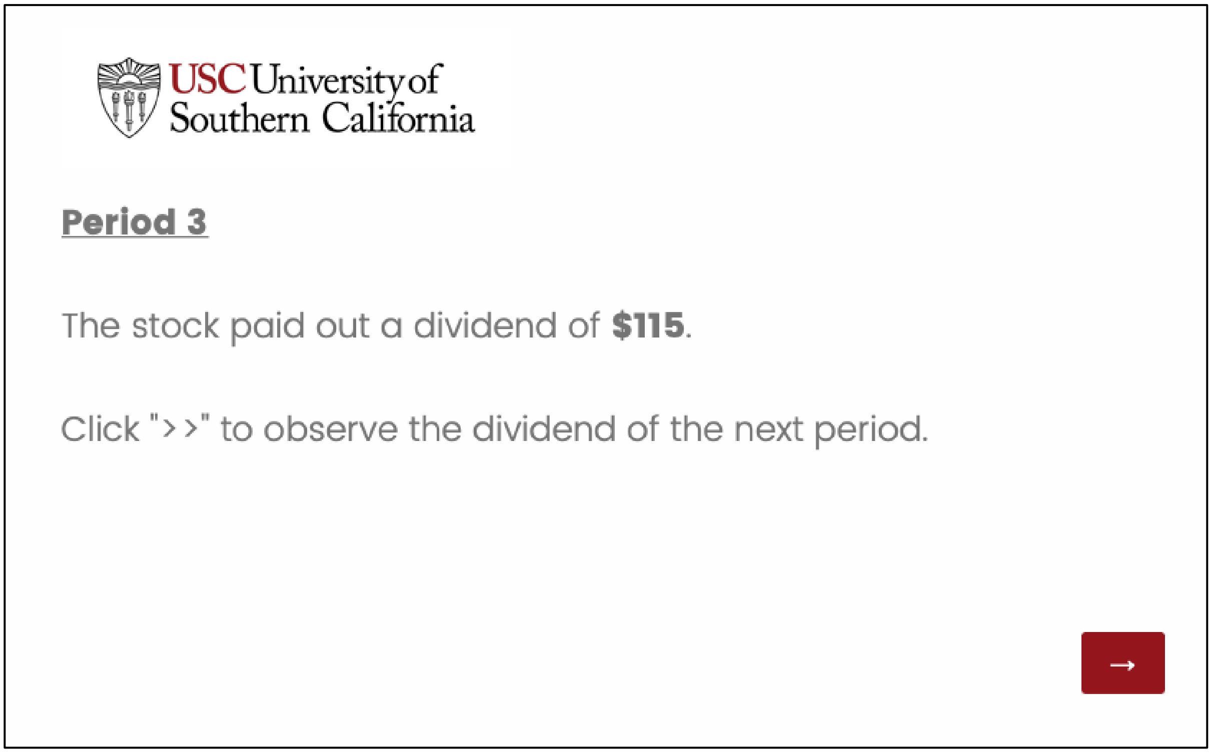

After observing the dividend realizations over the course of several periods, subjects were asked to answer two questions. While answering these questions, they received an overview of the full history of dividend realizations:

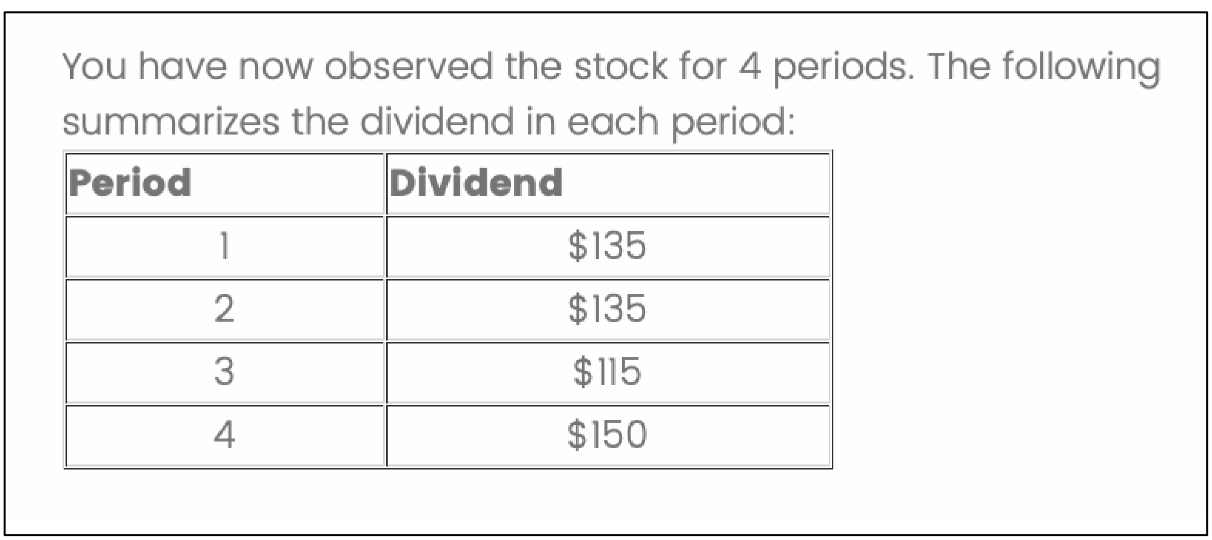
You have now observed the stock for 4 periods. The following summarizes the dividend in each period:

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

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%.

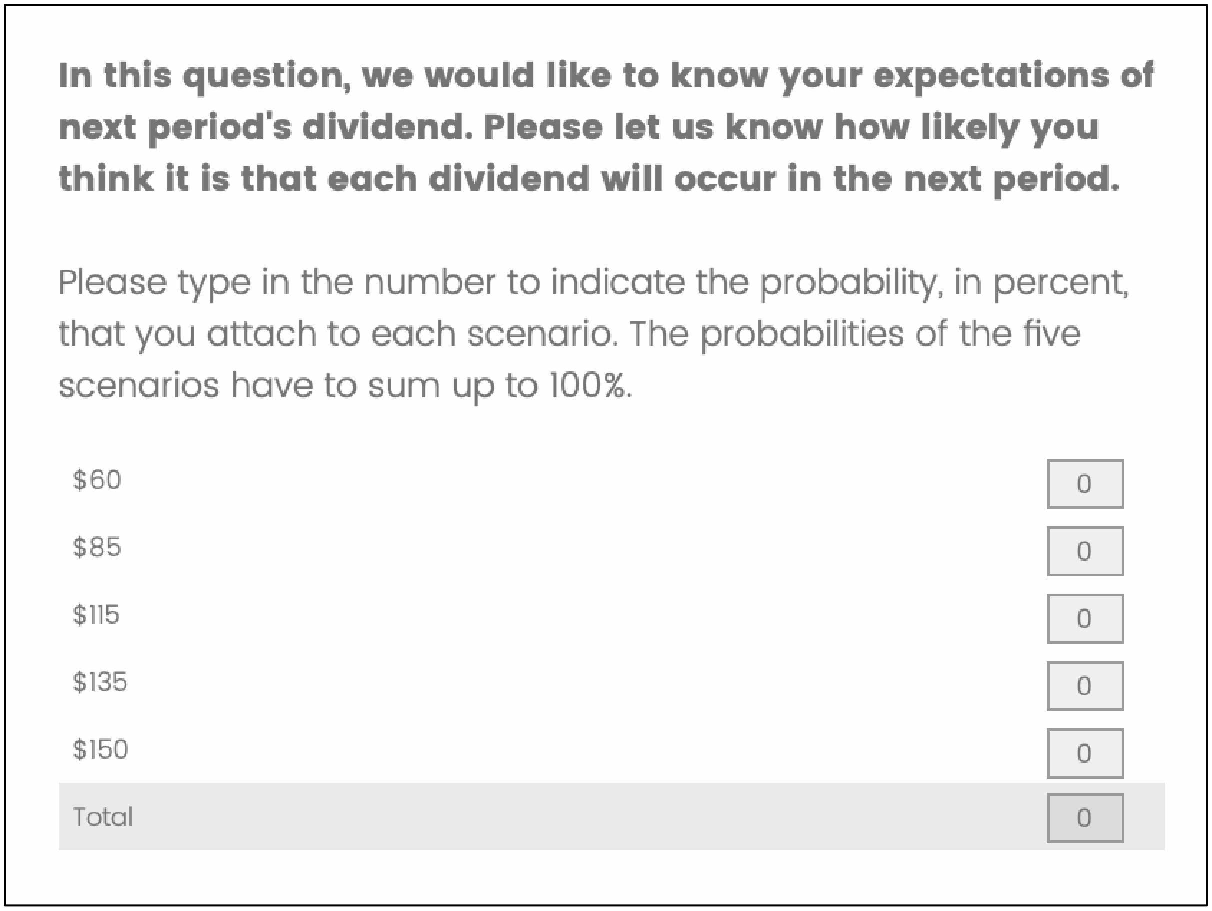

In this question, we would like to know your expectations of next period's dividend. Please let us know how likely you think it is that each dividend will occur in the next period.

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

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

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. Below are screenshots showing how the slider appeared before and after initiation: