Conservative Holdings, Aggressive Trades: Learning, Equilibrium Flows, and Risk Premia*

Thomas Dangl† **Lorenzo Garlappi**‡ **Alex Weissensteiner**§

July 14, 2022

*We thank Ella Dias Saraiva-Patelli, Alessandro Melone, Raman Uppal, Youchang Wu, Goufu Zhou, and participants at the Computational and Financial Econometrics conference in London, and at Seminars at University of Liechtenstein and WU Wien for helpful comments.

†Vienna University of Technology, Theresianumgasse 27, A-1040 Vienna, Austria. E-mail: thomas.dangl@tuwien.ac.at

‡Sauder School of Business at the University of British Columbia, 2053 Main Mall, Vancouver, BC V6T 1Z2, Canada. E-mail: lorenzo.garlappi@sauder.ubc.ca

§Free University of Bozen-Bolzano School of Economics and Management. Universitätsplatz 1 - Piazza Università 1, 39100 Bozen-Bolzano. E-mail: Alex.Weissensteiner@unibz.it

Conservative Holdings, Aggressive Trades: Learning, Equilibrium Flows, and Risk Premia

Abstract

We propose an equilibrium asset pricing model in which agents learn about fundamentals and differ in their preference for robustness. We first show that, when agents are averse to parameter uncertainty, learning about the volatility of fundamentals has a first-order effect on portfolio flows: contrary to intuition, uncertainty-averse agents *increase* their risky asset holdings in periods of high uncertainty, despite holding conservative portfolios. We then show that subjective risk premia increase following both unexpected good and bad news, implying that "betting-against-beta" strategies are less profitable when uncertainty is high. These predictions are consistent with observed portfolio flows of retail and institutional investors around dividend surprises and with recent empirical evidence linking uncertainty and the beta anomaly. Our model highlights that heterogeneity of preferences and learning about economic uncertainty are key channels for understanding the equilibrium dynamics of portfolio holdings and risk premia following news about economic outcomes.

JEL Classification Codes: G11, G12

Keywords: Ambiguity, uncertainty, learning, portfolio flows, equilibrium asset prices, heterogeneous agents

1 Introduction

Periods of high uncertainty, such as those following unexpected corporate and macro announcements, are frequently associated with a flow of risky assets from institutional to individual investors' portfolios. The existing literature typically attributes such flows to either investors' limited attention or portfolio constraints.[1] In this paper we propose an alternative explanation that emphasizes the equilibrium interaction between heterogeneous agents who, upon the arrival of new information, (i) learn about the underlying parameters of the economy and (ii) are sensitive to parameter uncertainty. We show that learning and aversion to parameter uncertainty are necessary channels to explain the equilibrium dynamics of portfolio flows and asset prices around economic announcements.

Recent evidence (see, e.g., Nagel and Xu, 2022) emphasizes that subjective, ex-ante, risk premia inferred from return expectations of investors for a variety of asset classes are much less counter-cyclical than objective risk-premia inferred from in-sample predictive regressions. Such a discrepancy between subjective and objective risk premia can be reconciled through representative-agent asset pricing models in which subjective expectations are time varying as in settings with constant learning (Collin-Dufresne, Johannes, and Lochstoer, 2016a; Nagel and Xu, 2021). While learning about uncertain parameters in representative-agent models introduces important dynamics that are relevant to explain asset pricing puzzles, these models are not designed to study portfolio flows. In this paper, we consider a general equilibrium heterogeneous-agents endowment economy in which different types of agents learn about both the mean and the volatility of the endowment process and differ in their aversion to parameter uncertainty. It is common in the literature to either (i) assume that dividend volatility is known as it can be perfectly estimated from high-frequency data or (ii) treat volatility as known invoking studies in which learning about volatility has a second-order effect in equilibrium, e.g., Collin-Dufresne, Johannes, and Lochstoer (2016b). In reality, however, information reaches market participants in a lumpy fashion, such as during FOMC communication events or corporate earning announcements, and agents have to learn about volatility. We show that, in the presence of aversion to uncertainty, learning about volatility has a first-order effect on portfolio decisions and gives rise to novel dynamics in portfolio flows and asset prices that are consistent with empirical observation.

[1] See, among others, Frazzini and Lamont (2007), Barber and Odean (2008) and Hirshleifer, Myers, Myers, and Teoh (2008) Kaniel, Saar, and Titman (2008), Kaniel, Liu, Saar, and Titman (2012).

Our model predicts that, in equilibrium, ambiguity averse agents, while more conservative in their holdings, are more aggressive in their trades. Specifically, they *increase* their position in the risky asset after observing large positive or negative dividend realizations. Furthermore, the agents' equilibrium flow adjustments to the arrival of information in our model implies that larger dividend realizations are associated with higher expected market risk premia. These results highlight the importance of considering equilibrium forces when interpreting empirical evidence.

We first illustrate the main mechanism in a stylized heterogeneous agent model that we can solve analytically. In this setting, we show that the equilibrium interaction between ambiguity-averse and ambiguity-neutral agents generates risk premia that depend linearly on both the dividend variance *and* the standard error of the expected dividend. This property is common when agents have preferences that exhibit first-order risk aversion, see, e.g., Segal and Spivak (1990). As a consequence of learning, both good and bad news result in an increase in agents' volatility estimate. When agents have heterogeneous preferences for robustness, these belief updates result in equilibrium in expected risk premia that are "too low" (i.e., prices too high) for ambiguity-neutral agents and "too high" (i.e., prices too low) for ambiguity averse agents to justify their existing portfolio holdings. This difference in valuation generates therefore gains from trade in which ambiguity averse (neutral) agents increase (decrease) their position in the risky asset in times of high uncertainty, that is, following both positive and negative surprises. In contrast, if no agent exhibits preference for robustness, the equilibrium risk premium is proportional to the dividend variance and a standard "no-trade" result emerges, where surprises do not generate equilibrium flows.

We then solve a fully-fledged multi-period general equilibrium endowment economy where agents learn about the mean and the variance of the endowment, and differ in their degree of risk aversion and preference for robustness. To gain tractability, we assume that agents in this model learn with "fading memory" as in Nagel and Xu (2021). We numerically solve such a model by extending the incomplete-market backward approach of Dumas and Lyasoff (2012) in order to account for learning and heterogeneity in the preference for robustness. Analysis of the equilibrium confirms the intuition of the stylized setting: ambiguity averse agents are conservative in their holdings but aggressive in their trade on new information; and, subjective risk premia are higher following both positive and negative surprises. Furthermore, we show that in equilibrium (i) the risk-free rate increases following negative surprises and decreases following positive surprises; and (ii) the aggressive trades of ambiguity averse agents give rise to patterns of over-reaction to dividend news in their equilibrium consumption.

We empirically investigate our model predictions using data on corporate ownership as well as changes in institutional ownership of individual public firms from F-13 filings. We find that exceptionally bad as well as exceptionally good signals of corporate profitability are associated with low or even negative changes in institutional ownership. In contrast, neutral signal realizations, indicating lack of surprise, are associated with an increase in institutional ownership. These findings are consistent with the predictions of our equilibrium model if we assume that, as suggested by the "competence hypothesis" (see, e.g., Heath and Tversky, 1991), retail investors are more averse to uncertainty than institutions.[2] Furthermore, we also find that, consistent with our model, the market risk premium is higher following negative as well as positive surprises. This finding is consistent with Nagel and Xu (2022) who report that subjective risk premia increase with the subjective estimate of variance.

In summary, within the context of our model (i) heterogeneous preferences for robustness; (ii) learning about the variance of the endowment process; and (iii) market clearing are necessary conditions to generate equilibrium portfolio flows and risk premia consistent with those observed in the data. In fact, a model without heterogeneity in preferences towards robustness cannot explain observed flows in response to news. Similarly, a model in which the variance of the dividend process is known does not generate sensitivity of portfolio flows to news. Finally, in a partial equilibrium model that ignores the price effects of portfolio rebalancing, as uncertainty increases, agents with preference for robustness trade into *more* conservative portfolios, contrary to the evidence we document in the data.

Our work relates to three strands of literature. First, we contribute to the literature on asset pricing with heterogeneous agents.[3] We differ from the work in this literature by considering learning and agents' preference for robustness emerging from their aversion to parameter uncertainty. Chapman and Polkovnichenko (2009) study asset pricing in two-date economies with heterogeneous agents endowed with non-expected utility preferences. We focus on one form of deviation from expected utility, namely preference for robustness, and generalize their results to the case of learning about the mean and the variance of the

[2]The "competence hypothesis" states that agents are generally ambiguity-averse toward tasks for which they do not feel competent. Li, Tiwari, and Tong (2017) provide support to the assumption that retail investors have a stronger desire for robustness.

[3]This literature is too vast to be reviewed here. Key contributions, among many others, are Mankiw (1986), Dumas (1989), Constantinides and Duffie (1996), Dumas, Kurshev, and Uppal (2009), Bhamra and Uppal (2014), and Gârleanu and Panageas (2015). Panageas (2020) provides an excellent review of the literature.

endowment process, and to multi-period economies.[4,5] Buss, Uppal, and Vilkov (2021) study the dynamics of asset demand in a multiperiod general equilibrium model in which agents are heterogeneous in their confidence about the assets' return dynamics. They show that heterogeneous beliefs lead to asset demand curves that are steeper than with homogeneous beliefs. Unlike Buss, Uppal, and Vilkov (2021), agents in our model differ in their attitude towards robustness and learn about both the mean and the variance of the dividend process. Because of agents' preference for robustness, learning about variance has a first-order effect on both equilibrium flows and asset prices in our model. These effects are instead negligible in a model where agents are ambiguity neutral.

Second, we contribute to the literature that studies asset prices under parameter uncertainty and learning.[6] In our model, agents update their priors about the moments of the endowment process by observing its realizations over time. If the endowment process were observable in continuous time, its variance could be known with certainty. In reality, information reaches market participants through discrete events, such as FOMC communications, earning announcements, macro news, etc.[7]

Learning in this setting implies that the estimated variance is time-varying. We show that time variation in the estimated variance has significant qualitative implications for equilibrium flows and asset prices that would be absent in a model where the variance is a known constant. An alternative approach would be to exogenously assume stochastic volatility in the endowment process (see, e.g. Bansal and Yaron, 2004). However, we show that in a model with unknown mean but stochastic and observable volatility, new observations always increase the precision of the mean and lead to negligible effects on portfolio flows. Therefore, we view learning about the volatility of a simple i.i.d. endowment process as a more

[4]Similar to our setup, Easley and O'Hara (2009) model investors with a desire for robustness with respect to ambiguity in both the dividend mean and variance. In our model, learning ties the ambiguity in the dividend mean to the variance of the dividend distribution and helps rationalize portfolio flows in reaction to new information. Cao, Wang, and Zhang (2005) use a similar model with heterogeneous uncertainty-averse investors but no learning to show that limited asset market participation can arise endogenously in the presence of model uncertainty.

[5]Illeditsch, Ganguli, and Condie (2021) analyze learning under ambiguity about the link between information and asset payoffs and show that this leads to underrection to news. Ilut and Schneider (2022) provide a comprehensive survey of modelling uncertainty as ambiguity.

[6]Among others, key contributions are David (1997), Veronesi (1999), Pástor (2000), Barberis (2000), Xia (2001) and Leippold, Trojani, and Vanini (2008). Pástor and Veronesi (2009) provide an extensive overview of learning in financial markets.

[7]See the large literature on the announcement premium, e.g., Savor and Wilson (2016), Ai and Bansal (2018), and many others.

realistic and parsimonious way to model time variation in volatility while allowing objective and subjective expectations to differ (see, e.g. Nagel and Xu, 2021).

Third, our work is related to the large literature that studies the trading behavior of institutional and retail investors. Ample evidence indicates that retail investors act as liquidity providers who meet institutional investors' demand for immediacy.[8] Consistent with this view, we document that institutional investors tend to reduce their share in corporate ownership when indicators of future corporate profitability are exceptionally bad. Although retail investors might be less sophisticated, they face lower agency costs and less liquidity constraints than their institutional counterparts. This advantage allows them to act as market makers, especially during times of financial turmoil when liquidity is a scarce resource. Surprisingly, and less discussed in this literature, institutional investors significantly reduce their share in corporate ownership after exceptionally *positive* signals as well. The finding that individual investors increase their holdings in the risky asset after bad *and* good surprises is rationalized in the literature by appealing to the "attention grabbing" hypothesis, which assumes that individual investors have limited attention and rarely short sell.[9] We provide an alternative explanation to the attention grabbing hypothesis, by highlighting the important role of preference for robustness and learning.

The rest of the paper proceeds as follows. In Section 2 we provide intuition in a simple one-period equilibrium model which is analytically tractable. Section 3 presents a general equilibrium heterogeneous-agent model with learning. Section 4 contains our empirical analysis. Section 5 concludes. Appendix A derives the properties of the agents' iso-portfolio curves used to develop intuition in the stylized model of Section 2. Appendix B describes Bayesian updating with fading memory and Appendix C contains details of the numerical algorithm we develop to solve the model of Section 3.

2 A one-period model

In this section we develop a simple one-period equilibrium model with learning and heterogeneity in the preference for robustness that we can solve analytically. We use this stylized model to illustrate the effect of learning about dividend volatility when agents are ambigu-

[8] See e.g. Kaniel, Saar, and Titman (2008), Barrot, Kaniel, and Sraer (2016), Glossner, Matos, Ramelli, and Wagner (2020), and Pástor and Vorsatz (2020).

[9] See e.g. Frazzini and Lamont (2007), Barber and Odean (2008), Hirshleifer, Myers, Myers, and Teoh (2008),Berkman, Koch, Tuttle, and Zhang (2012), and Barber, Huang, Odean, and Schwarz (2021).

ity averse (Section 2.2); we compare these effects to those obtained from a model in which volatility is known but stochastic (Section 2.3); and discuss the cyclicality of the equilibrium subjective risk premium (Section 2.4).

2.1 Setup

There are two dates, $t = 0, 1$, and a single "tree" producing a perishable dividend $\tilde{D}$ at time $t = 1$. Agents can trade in claims over the dividend tree (the risky asset) at a price p and lend to each other (the riskless asset). Since consumption occurs only at $t = 1$, the riskless rate in the economy is undetermined and, without loss of generality, set to zero.

Beliefs. We assume that the dividend $\tilde{D}$ is normally distributed with unknown mean μ and variance σ^2. At time $t = 0$ agents have subjective beliefs m and v^2 about μ and σ^2 that result from having observed a time series of n past dividend realizations. Thus, agents believe that the $t = 1$ dividend $\tilde{D}$ is t-distributed with mean m and variance v^2 and $n - 1$ degrees of freedom. For ease of exposition, we assume a large enough n, so that the subjective distribution of dividend is approximately normal.

Preferences. The economy is populated by two types of agents, $i = A, B$, both having CARA utility $u_i(c) = -1/\gamma_i e^{-\gamma_i c}$, with absolute risk aversion $\gamma_i > 0$. Both agents have the same beliefs (m, v^2) but differ in their attitude towards uncertainty in these moments' estimates. Type-B agents are Bayesians. Therefore, they (i) use their best estimate of the mean m as the "perceived" dividend mean, that is, $\hat{\mu}_B = m$, and (ii) account for parameter uncertainty by inflating their best estimate of the variance v^2 by the mean standard error $s = v/\sqrt{n}$, that is, their perceived dividend variance is $\hat{\sigma}^2 = v^2 + s^2 = v^2(1 + \frac{1}{n})$.[10]

In contrast, type-A agents are "ambiguity averse" and dislike the uncertainty in the parameter estimates m and v^2. They guard against their aversion to parameter uncertainty by choosing portfolios that are robust to "worst-case" scenarios. This implies choosing the perceived dividend mean and variance conservatively within a confidence interval for the estimates m and v^2. The size of these confidence interval captures agents' aversion to uncertainty. For illustration purposes, in this section we assume that the agent is averse only

[10] Because the standard error of the mean declines monotonically in the number of observations, n, the true mean becomes eventually known. In the next section, we allow for parameter uncertainty to persists over time by assuming that agents learn with "fading memory" as in Nagel and Xu (2021).

to uncertainty about the mean estimate m and consider the general case in Appendix B.2.[11] Therefore, like the Bayesians, ambiguity-averse agents account for parameter uncertainty by inflating their variance estimate v^2, that is, $\hat{\sigma}^2 = v^2+s^2$. In addition, they select the perceived mean $\hat{\mu}_A$ in a conservative way within a confidence interval $\hat{\mu}_A \in \mathcal{M} \equiv [m-\kappa s, m+\kappa s]$, whose size depends on their aversion to uncertainty κ. In classical statistics, κ would represent, for example, the quantile of a distribution, e.g., $\kappa = 1.96$ for a 95% confidence interval of a normal distribution. When $\kappa = 0$ and $\gamma_A = \gamma_B$, agent A and B are identical. Therefore, the parameter κ parsimoniously captures the heterogeneity in attitude towards parameter uncertainty between the agents in our economy.

Portfolios. We denote by θ_i agent's $i = A, B$ number of shares of the risky asset. The portfolio problem of type-B agents is standard, that is,

$$\max_{\theta_B} \mathbb{E}\left[-\frac{1}{\gamma_B} e^{-\gamma_B \tilde{W}_B}\right], \tag{1}$$

subject to

$$\tilde{W}_B \quad = \quad \theta_B \tilde{D} + (W_B - \theta_B p), \tag{2}$$

where W_B denotes the value of B's initial endowment.[12] If n is large enough, $\tilde{D} \sim \mathcal{N}(\hat{\mu}_B, \hat{\sigma}^2)$, with $\hat{\mu}_B = m$ and $\hat{\sigma}^2 = v^2(1 + \frac{1}{n})$. Therefore, agent B's demand of the risky asset is

$$\theta_B = \frac{m - p}{\gamma_B \hat{\sigma}^2}. \tag{3}$$

[11] Easley and O'Hara (2009) show that when ambiguity aversion is modeled as in the max-min setting of Gilboa and Schmeidler (1989), the ambiguity-averse agent always elects the highest possible return variance when constructing optimal portfolios. Therefore, the portfolio choice problem with ambiguity about both μ and σ^2 reduces to a problem with ambiguity only about μ, where σ^2 is fixed at the maximum over the set of its possible values.

[12] With CARA preferences and absent a consumption decision at $t = 0$, i.e., with exogenous riskfree rate, the demand for the risky investment is independent of the initial endowment. Hence, the relative share of the types of agents only depends on the relation between their risk tolerance, $\frac{1}{\gamma_i}$.

In contrast, type-A agents maximize expected utility by selecting the "worst case" mean $\hat{\mu}_A$ from their estimated confidence interval.[13] Specifically, type-A agents solve

$$\max_{\theta_A} \min_{\hat{\mu}_A \in \mathcal{M}} \mathbb{E}\left[-\frac{1}{\gamma_A} e^{-\gamma_A \tilde{W}_A}\right], \quad \text{with } \mathcal{M} \equiv [m - \kappa s, m + \kappa s], \tag{4}$$

subject to

$$\tilde{W}_A = \theta_A \tilde{D} + (W_A - \theta_A p), \tag{5}$$

with W_A the value of A's initial endowment. The desire for robustness of types-A agents implies that, when forming their consumption and portfolio decisions, type-A agents select the mean $\hat{\mu}_A$ within the confidence interval $\mathcal{M}$ that delivers the worst possible expected utility.[14]

The solution of the portfolio problem (4)–(5) is

$$\theta_A = \begin{cases} \frac{m-\kappa s-p}{\gamma_A \hat{\sigma}^2} & \text{if } p < m - \kappa s, \\ 0 & \text{if } m - \kappa s \leq p \leq m + \kappa s. \\ \frac{m+\kappa s-p}{\gamma_A \hat{\sigma}^2} & \text{if } p > m + \kappa s \end{cases} \tag{6}$$

When $p < m - \kappa s$, agents A have positive demand for the risky asset, and they invest as if they would use a different (equivalent) probability measure on the dividend states. That is, type-A agents invest as if their estimate of the expected dividend is adjusted downwards by a quantity κs that depends on their degree of ambiguity aversion κ and the amount of uncertainty in the mean estimate m, captured by the standard error s. When $p \in (m - \kappa s, m + \kappa s)$ type-A agents do not hold the risky asset and for $p > m + \kappa s$, they hold a short position.

Equilibrium. By imposing market clearing, $\theta_A + \theta_B = 1$, we obtain the equilibrium ex-dividend price for the risky asset

$$p^* = m - \lambda, \tag{7}$$

[13]For simplicity, in our analysis we rely on the "max-min" implementation of the Gilboa and Schmeidler (1989) model, as in Garlappi, Uppal, and Wang (2007). Alternative and less extreme versions of this approach are possible, such as models with "variational preferences" as in Hansen and Sargent (2001), in which the desire for robustness can be captured by a "penalty" for deviations from the posterior (m_t, v_t^2), see, for example Anderson, Hansen, and Sargent (2000) and Hansen and Sargent (2008). The main qualitative implications are however unaltered: aversion to parameter uncertainty leads to a direct adjustment in the perceived mean estimate before forming portfolios.

[14]See Bewley (2011) for a discussion of how confidence intervals obtained from classical statistics are related to Knightian uncertainty.

where λ denotes the subjective risk premium given by

$$\lambda = \begin{cases} \frac{\gamma_B}{\gamma_A+\gamma_B}\frac{\kappa}{\sqrt{n}}v + \gamma_0\left(\frac{n+1}{n}\right)v^2 & \text{if } \kappa \leq \kappa^*, \\ \gamma_B\left(\frac{n+1}{n}\right)v^2 & \text{if } \kappa > \kappa^*. \end{cases} \quad \text{with} \quad \kappa^* \equiv \gamma_B\left(\frac{n+1}{\sqrt{n}}\right)v, \tag{8}$$

where $\gamma_0 \equiv \left(\gamma_A^{-1} + \gamma_B^{-1}\right)^{-1}$ denotes the aggregated absolute risk aversion in the economy.

The demand for the risky asset in equations (3) and (6) imply that either both agent participate in the market or only agent B participates. The expression of the risk premium in equation (8) shows that A participates for level of ambiguity aversion smaller than the threshold κ^*. Intuitively, A's demand is higher with less aversion to ambiguity. From the expression of the participation threshold κ^* we note that equilibrium participation is more likely the higher is (i) B's risk aversion γ_B ; (ii) the estimated dividend volatility v; and (iii) the number of past dividend observations n.

The expression for the equilibrium subjective risk premium in equation (8) shows that when both agents participate in the market, i.e., $\kappa < \kappa^*$, the equilibrium risk premium is linear-quadratic in the dividend volatility v. The linear term in the expression of λ appears because the preferences of type-A agents exhibit "first-order" risk aversion, (see, e.g., Segal and Spivak, 1990). Intuitively, unlike B agents who are locally risk-neutral, A agents are locally risk-averse and demand a compensation for holding a vanishing amount of risk.[15] Note that, from equation (8), the subjective risk premium λ depends only on the dividend volatility v and variance v^2 and it will therefore be constant in an economy in which the dividend variance is a known constant.

Substituting the equilibrium price p^* in the agents' demand functions (3) and (6), we obtain the following expressions for the equilibrium portfolio holdings[16]

$$\theta_A = \max\left\{\frac{\gamma_0}{\gamma_A} - \frac{1}{\gamma_A+\gamma_B}\left(\frac{\sqrt{n}}{n+1}\right)\frac{\kappa}{v}, 0\right\}, \tag{9}$$

$$\theta_B = \min\left\{\frac{\gamma_0}{\gamma_B} + \frac{1}{\gamma_A+\gamma_B}\left(\frac{\sqrt{n}}{n+1}\right)\frac{\kappa}{v}, 1\right\}. \tag{10}$$

[**Figure 1 about here.**]

[15] Aversion towards uncertainty in σ^2 does not generate fist-order risk aversion and has only quantitative, not qualitative, effects. See Appendix B.2 for a formal argument.

[16] Note, that the max and min in equations (9) and (10) do not originate from a short-sale constraint but from agent A's no participation, as shown in the demand equation (6).

The left panel of Figure 1 illustrates the equilibrium portfolio weights θ_A and θ_B from equations (9)–(10) as a function of the volatility estimate v. The figure shows that if $\kappa < \kappa^*$ or, equivalently, $v > v^* \equiv \frac{\sqrt{n}}{n+1}\frac{\kappa}{\gamma_B}$ (the vertical dashed line), both agents hold the risky asset in equilibrium. Furthermore, when both agents participate, A's risky asset demand is increasing in the dividend volatility v while B's demand is decreasing. As $v \to \infty$, the portfolio holdings converge asymptotically to the constant weights $\theta_A = \frac{\gamma_0}{\gamma_A}$ and $\theta_B = \frac{\gamma_0}{\gamma_B}$.

The right panel of Figure 1 provides an intuition for the structure of the equilibrium holdings in equations (9)–(10). The dotted curves in the figures represent "iso-portfolio" curves for both agents, that is, the combination of volatility v and risk premium λ associated with the same portfolio holdings. The intersection of complementary iso-portfolio curves of the two agents, i.e., curves with weights that clear the market, $\theta_B + \theta_A = 1$, identify the equilibrium risk premium, represented by the solid black line. Because of first-order risk aversion, agents A participate in the market only if the risk premium exceeds the hurdle $\lambda(v^*) = \frac{\kappa^2}{(n+1)\gamma_B} = \frac{\kappa}{\sqrt{n}}v^*$. The red-shaded area in Figure 1 shows the range of risk premia that are too low for A to hold the risky asset. The $\theta_A = 0$ iso-portfolio line (red-dashed) intersects the curve of the equilibrium risk premium (black line) at $v = v^*$ where $\lambda(v^*) = \frac{\kappa}{\sqrt{n}}v^*$. For a dividend volatility $v \leq v^*$, or, equivalently, $\kappa \geq \kappa^*$, the risk premium that agent B requires for holding 100% of the risky asset is too low for agents A to participate in the market. Therefore, since agent B is the only one holding the risky asset, the risk premium's dependence on v coincides with the $\theta_B = 100\%$ iso-curve (upmost blue dashed line).

For values of $v > v^*$, or, equivalently, $\kappa < \kappa^*$, both agents participate in equilibrium. As we show in Appendix A, in any equilibrium in which agents A participate, their iso-portfolio lines in Figure 1 are always flatter than those of agents B. This happens because, when agents exhibit first-order risk aversion, the equilibrium risk premium λ consists of two parts: (i) compensation for first-order risk aversion to induce participation (v-term in equation (8)) and (ii) compensation for risk aversion (v^2-term in equation (8)). Because of this dual function of the risk premium, agents A hold fewer risky assets than B and demand a relatively lower compensation for accepting an additional marginal unit of volatility. This implies that the marginal rate of substitution between required risk premium and dividend volatility is strictly higher for B than for A. Intuitively, starting from an equilibrium in which both A and B participate, in order not to change their portfolios, agents A require relatively less compensation than B for an additional unit of volatility. Hence the equilibrium risk premium is perceived to be "too high" (price too low) by A and "too low" (price too high) by B, and a gain from trade emerges: B is willing to sell and A is willing to buy. In sum, because of first-order risk aversion, in equilibrium agents A hold "conservative" portfolio but trade "aggressively" by increasing the holdings of the risky asset following an increase in volatility.

2.2 Learning and equilibrium flows

As the analysis of the previous section illustrates, in an equilibrium when both types of agents participate, portfolio holdings depend on the estimated dividend volatility, v and, from equations (9)–(10), changes in v induces portfolio flows between the Bayesian and the ambiguity averse investors. This fact suggests that learning about volatility, by making the estimate v time varying, plays a crucial role in the determination of equilibrium flows. Here we develop intuition about the effect of learning in the context of the simple model of the previous section. The general model of Section 3 explores this channel thoroughly.

Suppose agents form their $t = 0$ beliefs (m, v^2) from a set of n historical dividend observations, resulting in the equilibrium portfolio weights θ_A and θ_B shown in equations (9)–(10). At time $t = 0$, agents observe a dividend signal D'. Because agents have CARA preferences and there is no consumption at $t = 0$, the signal D' has no wealth effects and therefore does not induce agents to rebalance their portfolio. The only effect of the signal D' is to reveal information about future dividends. After observing this signal, investors update their belief to (m', v'^2) resulting in new equilibrium holdings, θ'_A and θ'_B. We interpret the rebalancing $\theta'_A - \theta_A = \theta_B - \theta'_B$ as the flow induced by the arrival of new information. The difference $e = D' - m$ between the new observation and the agents' prior belief m represents the dividend "surprise". Standard results from statistical inference theory, see, e.g., Greene (2020) imply that the posterior mean, m', is linear in the surprise e and the posterior variance, v'^2, is linear in the squared surprise e^2 , that is,

$$m' = \frac{1}{n'} \sum_{t=1}^{n'} D_t = m + \frac{1}{n'} e, \tag{11}$$

$$v'^2 = \frac{1}{n'-1} \sum_{t=1}^{n'} \left(D_t - m' \right)^2 = v^2 + \frac{1}{n'} \left(e^2 - \frac{n'}{n'-1} v^2 \right), \tag{12}$$

$$s' = \frac{v'}{\sqrt{n'}}. \tag{13}$$

with $n' = n + 1$ the number of observations, and $\frac{n'}{n'-1}v^2 = v^2 + s^2$ the prior estimate of the total variance.

The portfolio holdings (9)–(10) highlight that only a change in volatility can generate trade among agents in equilibrium. Because of CARA preferences, these equilibrium portfolios do not depend on the agents' beliefs about the dividend mean, m. Hence, if the dividend variance is known there would not be equilibrium flows in this model. Equilibrium flows can emerge only if v varies over time, as in the case of a model in which agents learn about the variance of the dividend process. Larger dividend "surprises" e lead to an increase in the updated variance v'^2 and hence in the standard error $s' = v'/\sqrt{n'}$ of the estimated mean. Therefore, from equations (9)–(10), large surprises imply an *increase* in the equilibrium holdings θ_A and a *decrease* in the equilibrium holdings θ_B.

[**Figure 2 about here.**]

Figure 2 shows equilibrium portfolio holdings (left panels) and subjective risk premia (right panels) for low (top panels) and high (bottom panels) levels of ambiguity aversion, as a function of the dividend surprise $e = D' - m$. Unlike Figure 1, which shows the same quantities as a function of the volatility estimate v, Figure 2 emphasizes the effect of both positive and negative surprises. A large value of the volatility in Figure 1 can be associated with either a large positive or negative dividend surprise. Linking portfolio flows and risk premia to dividend surprises allows us to relate more directly to the empirical analysis of Section 4.

The left panels of Figure 2 illustrate the equilibrium portfolio holdings of agents B from equation (10) as a function of the dividend surprise e. The different curves correspond to different levels of type-A agents' preference for robustness parameter, κ. Larger values of κ imply stronger ambiguity aversion, more conservative portfolios for agents A, and hence, by market clearing, larger risky asset holdings by B agents. The inverted U-shape of the equilibrium portfolio of agents B indicates that larger dividend surprises reduce the holdings θ_B of the Bayesian investor and increase those of the ambiguity-averse agent A. The black line in the top-left panel shows that with no preference for robustness ($\kappa = 0$) there are no flows following dividend surprises. For intermediate values of κ there are portfolio flows between agents, with A-agents buying and B-agents selling. For higher values of κ (red line in the bottom-left panel), ambiguity averse agents do not participate and therefore there cannot be flows between agents in equilibrium.

The right panels of Figure 2 display the equilibrium risk premium λ as a function of dividend surprises. A larger dividend surprise e, positive or negative, is associated to larger values of the posterior variance v^2. It then follows directly from the mechanism illustrated in Figure 1, that larger surprises are associated with larger expected risk premia. The right panels of Figure 2 also show that, although a higher preference for robustness κ results in higher risk premia, conditional on participation of A-type agents, ambiguity aversion, per se, is not necessary to generate the U-shape relationship between dividend surprises and risk premia. In fact, such a pattern is present also for the case of $\kappa = 0$ (the black line in the top-right panel).

2.3 Learning about variance vs. stochastic volatility

One may argue that a model in which perceived variance is endogenously time-varying due to learning is observationally equivalent to a model without learning but with exogenously time-varying variance. Although both models exhibit time-variation in volatility, in a model with learning, a change in the estimated variance implies a change in the perceived information quality of *all* historically observed dividends. In contrast, in a model with stochastic volatility, a change in variance does not affect the quality of past information since parameters are known. In this section, we show that the difference in the source of time-variation in volatility has qualitative and quantitative implications for portfolio flows. Specifically, while known and stochastic volatility has only second-order effects on flows, learning about the volatility and its impact on the confidence interval of the mean has first-order effects in any equilibrium with ambiguity averse agents.

To illustrate this point, let us assume independent and identically normally distributed dividends D with unknown and constant mean μ and time-varying but observable variance σ_t^2. The best estimate of the mean m from a history of n observations is then (see, e.g., Chapter 9 in Greene, 2020)

$$m = \sum_{t=1}^{n} w_t D_t \quad \text{with} \quad w_t \equiv \frac{1/\sigma_t^2}{1/s^2} \quad \text{and} \quad \frac{1}{s^2} = \sum_{t=1}^{n} \frac{1}{\sigma_t^2}, \tag{14}$$

where the weight w_t represents the relative precision of each observation and s the standard error of the mean.[17]

The updated values of the mean and standard error after observing the new realized dividend D' and variance σ'^2 are

$$m' = (1 - w')m + w'D', \quad w' = \frac{\frac{1}{\sigma'^2}}{\frac{1}{s^2} + \frac{1}{\sigma'^2}} = \frac{s^2}{s^2 + \sigma'^2} \tag{15}$$

$$\frac{1}{s'^2} = \frac{1}{s^2} + \frac{1}{\sigma'^2}. \tag{16}$$

With stochastic but known variance, the updated standard error s' does not depend on the new dividend realization D'. Equation (16) implies that $s' \leq s$, i.e., new observations can only *reduce* the standard error of the mean.[18] Because the standard error controls the size of the confidence interval $\mathcal{M} = [m - \kappa s', m + \kappa s']$ characterizing ambiguity about the mean, in a model with stochastic volatility, a new dividend observation always reduces ambiguity. Observations in times of very high volatility σ' receive tiny weights w' in the updated mean m' and only marginally reduce the standard error s'.

In contrast, in our model where agents learn about an unknown variance, dividend surprises D' increase the estimated variance v'^2 whenever the squared surprise $e^2 = (D' - m)^2$ exceeds the prior estimate of the total variance $v^2 + s^2$, see equation (12). This increase in estimated variance directly implies a higher estimated standard error of the mean, $s' = v'/\sqrt{n'}$, see equation (13). Therefore, the new signal affects the quality of *all* historical dividends, and agents revise their confidence interval of the mean. Comparing the dynamics of the standard error of the mean in a model with stochastic volatility, equation (16), and learning about variance, equation (12), we conclude that a model of stochastic but know volatilty would imply negligible or inexistent equilibrium flows following dividend surprises, contrary to the empirical evidence.

[17] Because dividend realization are independent and identically distributed, the variance of m is given by

$$s^2 = \text{var}(m) = \sum_{t=1}^{n} w_t^2 \underbrace{\text{var}(D_t)}_{=\sigma_t^2} = \left(\frac{1}{\sum_{t=1}^{n} \frac{1}{\sigma_t^2}}\right)^2 \sum_{t=1}^{n} \left[\left(\frac{1}{\sigma_t^2}\right)^2 \sigma_t^2\right] = \frac{1}{\sum_{t=1}^{n} \frac{1}{\sigma_t^2}}.$$

[18] In a model of fading memory, as the one considered in Section 3, it is possible that new observations might lead to a slight increase in the standard error, when agents' memory fades sufficiently fast.

2.4 Countercyclical risk premia and predictability

The expression for the equilibrium risk premium in equation (8) and the updating in variance estimate (12) induced by learning implies that the subjective risk premium λ increases with the magnitude of the dividend surprises $e = D' - m$, regardless of their sign. This dependence gives rise to a form of predictability that investors in our model can detect in real time. As in Lewellen and Shanken (2002) and Nagel and Xu (2022), parameter uncertainty implies also a second form of predictability that, although detectable by an outside econometrician who knows the dividend process' true moments, cannot be exploited by investors.

To illustrate, consider first the case in which the true dividend mean is constant but unknown, and the variance is known by all investors. Following positive dividend realizations, investors' best estimate of the expected dividend is higher than the true mean and the stock will be "over-priced" relative to its fundamental value. Since the true mean is lower than investors' estimate, they will perceive negative returns after high prices. An econometrician looking at the data will find that high prices predict lower returns, that is, the objective expected risk premium is

$$\lambda^{\text{obj}} = \mu - p^* = \underbrace{\mu - m(e)}_{\text{unobservable}} + \underbrace{m(e) - p^*}_{\equiv \lambda^{\text{subj}}(e)}, \tag{17}$$

where $\lambda^{\text{subj}}(e)$ denotes the subjective risk-premium derived in equation (8), which is a constant if σ^2 is known. Following a positive dividend realization $e > 0$, agents' best estimate of the mean $m(e)$ increases and the objective expected risk premium λ^{obj} decreases, implying lower expected returns. However, unlike the econometrician, investors do not know the true dividend mean μ and therefore cannot exploit such a profitability. If variance is known, the subjective risk premium $\lambda^{\text{subj}}(e)$ is constant and there is no real-time predictability and hence no flows in equilibrium.

In contrast, when agents learn also about the variance, as Figure 2 illustrates, the equilibrium subjective risk-premium $\lambda^{\text{subj}}(e)$ increases with the size of the dividend surprise e, regardlesss of its sign. After large positive or negative surprises, agents expect returns to be higher.[19] Unlike the term $\mu - m(e)$, which cannot be known by the agents, the subjec-

[19] Nagel and Xu (2022) analyze CFO survey data and find that the subjective risk premium is positively related to subjective estimates of variance and that CFOs' subjective return expectations strongly depend on realized variance.

tive risk-premium $\lambda^{\text{subj}}(e)$ belongs to the agents' information sets. Therefore, in equilibrium, agents experience a time-varying risk premium.

The objective risk premium (17) is asymmetric with respect to new information. Because the mean estimate $m(e)$ increases with the surprise e and the subjective risk premmium $\lambda^{\text{subj}}(e)$ increases with the square surprise e^2, bad news ($e < 0$) result in a more pronounced increase in the objective risk premium (a larger drop in the share price) than "good news" ($e > 0$) of equal magnitude, which instead result in a more moderate decrease in the risk premium. This asymmetric reaction to new information is a direct consequence of learning and does not require additional behavioral assumptions, such as agents' over-reaction to bad news.

3 A multi-period model

The simplified model of the previous section is useful to develop intuition. The analytical solution of this model, however, comes at the cost of ignoring intertemporal consumption decisions, and hence is silent about the equilibrium risk-free rate and learning dynamics. In this section, we extend the setting to a general equilibrium model in which agents optimally choose their consumption and portfolio policy while learning about the moments of the dividend process. This general model confirms that the results and the intuition developed in the model of the previous section remain valid in a more general framework.

3.1 Setup

We consider a discrete-time endowment economy with heterogeneous agents living over a finite horizon T, with time indexed by $t = 0, \ldots, T$. There is a single consumption good which we take as the numéraire. The financial assets in this economy consist of two traded securities: (i) a short-lived riskless asset (the "bond") in zero-net supply and unit face value and (ii) a long-lived asset (the "stock") in unit supply. The stock generates a perishable dividend $D_t \sim \mathcal{N}(\mu, \sigma^2)$ at dates $t = 1, \ldots, T$. The dividend follows an independent and identically distributed (IID) law of motion

$$D_t = \mu + \sigma\epsilon_t, \quad t = 1, \ldots, T, \tag{18}$$

where $\{\epsilon_t\}$ is a series of standard normal shocks.[20] The dividend mean μ and its variance σ^2 are unknown to the agents and must be learned from observing the dividend realization D_t. We discuss the learning process in Section 3.2.

The economy is populated by two types of atomistic agents, $i = A, B$, each of equal mass.[21] At time $t = 0$ agents are endowed with common prior beliefs about the unknown mean and variance of the dividend process. At time $t = 1, \ldots, T$ agents observe the dividend D_t and update their beliefs. Based on these common posterior beliefs, they choose a level of consumption and a portfolio of stocks and bonds at time t. At time $t = T$ agents consume the liquidating value of their portfolio.

While both agents learn about the unknown parameters in the same way, they differ in how they use their posterior beliefs in forming consumption and portfolio decisions. As in Section 2, we assume that type-A agents are averse to parameter uncertainty and hence have a desire for robustness, while type-B agents are Bayesian and account for uncertainty only by inflating the perceived variance of the dividend process. We discuss the agents' preference specifications in more detail in Section 3.3.

3.2 Learning

Agents in our model learn about both the mean, μ, and the variance, σ^2, of the dividend process. Before observing the realized dividend D_t, agents have prior estimates (beliefs) m_{t-1} and v_{t-1}^2 of the unknown mean μ and variance σ^2, as well as an estimate of the standard error s_{t-1} of the mean estimate m_{t-1}. At time t, agents observe the dividend D_t and update their estimates.

A natural issue that emerges in a multi-period setting with constant but unknown parameters is that eventually investors learn the true parameters. In reality, parameter uncertainty is unlikely to disappear over time as the economy continuously evolves. To capture this fact in a tractable way, we assume that agents learn with "fading memory" as in Nagel and Xu (2021), implying that they put more emphasis on the recent history. This assumption pre-

[20]The dividend process in equation (18) implies that dividends are stationary in level, which is not realistic. We do so for consistency with the simple model of Section 2. In this sense, the analysis that follows can be thought of as applying to a detrended version of a model with deterministic growth, see, e.g., Epstein and Schneider (2008). Hence, without loss of generality, the methodology of this section can also be adapted to handle a model with stochastic growth.

[21]Both classes of agents are endowed with 50% of the long-lived risky asset and have equal risk aversion, $\gamma_A = \gamma_B$.

serves tractability by keeping agents in a realistic state of "perpetual learning" over time. Fading-memory investors' skepticism in the historical estimates implies an upward biased assessment of the perceived signal quality, and hence, in an increased sensitivity of the posterior beliefs to the more recent signal observation.[22]

In Appendix B we derive the asymptotic learning recurrence from standard theory of Bayesian filtering (see, e.g., West and Harrison, 2006, and their concept of "discounting" information) and show that this type of fading memory leads to "constant gain" learning (see, e.g., Evans and Honkapohja, 2009, 2012). The updating rules for posterior beliefs m_t and v_t^2 then simplify to the following constant-gain learning recurrences for both the estimated mean and variance of the dividend

$$m_t = aD_t + (1-a)m_{t-1} = m_{t-1} + a\underbrace{(D_t - m_{t-1})}_{\equiv e_t}, \tag{19}$$

$$v_t^2 = (1-a)v_{t-1}^2 + a(1-a)e_t^2, \tag{20}$$

$$s_t^2 = av_t^2, \tag{21}$$

where $e_t = D_t - m_{t-1}$ denotes the forecast error, or dividend surprise; s_t^2 denotes the estimation risk due to uncertainty about the mean dividend; and $a \in (0,1)$ is a constant Bayesian gain parameter that represents the perceived signal quality relative to the confidence in the prior and captures the responsiveness of the posterior beliefs to the arrival of new information.[23] The values m_t and v_t^2 in equations (19) and (20) are the investors' "best estimates" of the dividend mean μ and variance σ^2, respectively. Because the dividend mean is not known, the perceived total variance of dividend, $\hat{\sigma}_t^2$, consists of the sum of the "fundamental risk", v_t^2, and the "estimation risk", s_t^2, that is, $\hat{\sigma}_t^2 = v_t^2 + s_t^2 = (1+a)v_t^2$. One can think of the parameter a as of the ratio $1/n_{\text{eff}}$, with n_{eff} denoting the "effective" number of observations agents use to determine their beliefs, reflecting over-weighting of more recent observations.[24] In the standard Bayesian case of non-fading memory, the asymptotic gain for $n_{\text{eff}} \to \infty$ is $a = 0$. With an infinite number of observations, investors eventually acquire

[22] A similar learning mechanism that privileges recent history is at play also in the model of Collin-Dufresne, Johannes, and Lochstoer (2016a) and is documented empirically by Malmendier and Nagel (2011).

[23] The time-t posterior of μ has a Student's t-distribution with mean m_t, variance s_t^2, and $1/a$ degrees of freedom. The time-t posterior of σ^2 has a χ^2 distribution with $1/a$ degrees of freedom, see Appendix B for details.

[24] Bayesian forecasts that employ predictive regressions with time varying coefficients asymptotically lead to updating rules that are equivalent to learning under fading memory in equations (19)–(21), see, e.g., Koop and Korobilis (2013) or Dangl and Halling (2012).

perfect information about the moments (μ, σ) and, hence, do not update their beliefs upon new dividend realizations.

3.3 Preferences

Both types of agents update their beliefs according to the learning dynamics described in equations (19)–(21) with a common constant-gain learning parameter a. They differ, however, in the way they use such information when making their consumption and portfolio decisions. We assume that type-B agents are Bayesians while type-A agents are averse to parameter uncertainty and seek robustness when forming portfolios. To formalize this difference in investors' preferences, we assume that type-B agents use the posterior mean and variance estimates m_t and $\hat{\sigma}_t$ in forming their portfolio. In contrast, type-A agents consider a set of possible models, represented as confidence intervals around their best estimates of the dividend mean and variance, and form portfolios by considering a "worst case scenario". Specifically, type-A agents consider the following confidence intervals for the mean and variance of the dividend

$$\mathcal{M}_{\mu,t} \equiv [m_t - \kappa s_t, m_t + \kappa s_t], \quad \kappa > 0, \tag{22}$$

$$\mathcal{M}_{\sigma^2,t} \equiv \left[\underline{\ell} v_t^2, \bar{\ell} v_t^2\right], \quad 0 < \underline{\ell} < 1 < \bar{\ell}, \tag{23}$$

where κ, $\underline{\ell}$, and $\bar{\ell}$ are preference parameters capturing the agent aversion to uncertainty in the mean and variance, respectively, see Appendix B.1. The confidence intervals (22)–(23) represent the class of models that type-A agents deem feasible. Notice that learning affects both the location and the size of the confidence interval for the mean, $\mathcal{M}_{\mu,t}$. Larger unexpected surprise realization e_t result in larger updated variance v_t^2 and hence larger estimation risk s_t. As illustrated in the stylized model of Section 2, this process of learning about the variance is key to generate equilibrium flows and predictability in our model.

As in Section 2, we assume agents with time-separable CARA utility from consumption, c_t, that is, $u_i(c_t) = -\frac{1}{\gamma_i} e^{-\gamma_i c_t}$, for $i = A, B$ with γ_i denoting the coefficient of absolute risk aversion. Agents choose portfolio and consumption to maximize their expected discounted lifetime utility. Because of fading memory, the law of iterated expectations does not necessarily hold in equilibrium and therefore, as in Nagel and Xu (2021), we define agents' preferences in a recursive way, which guarantees that agents' consumption and portfolio policies are time-consistent. Formally,

$$\begin{aligned} U_A(c_t, c_{t+1}, \ldots, c_T) &= u_A(c_t) + \min_{(\hat{\mu}_{A,t}, \hat{\sigma}^2_{A,t}) \in \mathcal{M}_{\mu,t} \times \mathcal{M}_{\sigma^2,t}} \left\{ \beta \mathbb{E}_{A,t} \left[U_A(c_{t+1}, c_{t+2}, \ldots, c_T) \right] \right\}, \\ U_B(c_t, c_{t+1}, \ldots, c_T) &= u_B(c_t) + \mathbb{E}_{B,t} \left[\beta U_B(c_{t+1}, c_{t+2}, \ldots, c_T) \right], \\ U_i(c_T) &= u_i(c_T), \quad i = A, B, \end{aligned} \tag{24}$$

with β an impatience parameter, common to both agents. Conditional expectations reflect agents' beliefs about the parameters (μ, σ^2) of the data generating process.

Because of agents A's desire for robustness, the two types of agent de facto use different mean and variance estimates when computing their asset demand. These differences, however, are not due to heterogeneous beliefs, as both agents use the same learning technology, but are a consequence of their difference in attitude towards parameter uncertainty.

We denote by p_t^b the time-t price of the one-period bond with unit face value and by p_t^s the ex-dividend stock price at time t. At time $t = 0, \ldots, T$, each agent $i = A, B$ chooses the consumption $c_{i,t}$ and the portfolio of stocks and bonds, $(\theta^s_{i,t}, \theta^b_{i,t})$ to maximize their expected discounted lifetime utility given in (24) subject to the budget constraint

$$c_{i,t} + \theta^s_{i,t} p^s_t + \theta^b_{i,t} p^b_t = \theta^s_{i,t-1} \left(p^s_t + D_t \right) + \theta^b_{i,t-1} \cdot 1, \quad t = 0, \ldots, T, \tag{25}$$

with $\theta^s_{i,T} = \theta^b_{i,T} = 0$. We assume that at time $t = 0$ agents are endowed with $\theta^s_{i,-1}$ units of the stock and $\theta^b_{i,-1}$ units of the bond so that market clearing is satisfied, that is,

$$\theta^s_{A,-1} + \theta^s_{B,-1} = 1 \tag{26}$$

$$\theta^b_{A,-1} + \theta^b_{B,-1} = 0. \tag{27}$$

By construction, the ex-dividend stock price at the terminal date T is $p^s_T = 0$.

3.4 Equilibrium

A financial market equilibrium consists of price processes $\{p^b_t, p^s_t\}_{t=0}^T$, portfolios processes $\{\theta^b_{i,t}, \theta^s_{i,t}\}_{t=0}^T$, $i = A, B$, and consumption allocations $\{c_{A,t}, c_{B,t}\}_{t=0}^T$ such that (i) agents maximize their lifetime utility defined in equation (24) under the dynamic budget constraint (25); (ii) financial and goods market clear, that is, $\theta^s_{A,t} + \theta^s_{B,t} = 1$, $\theta^b_{A,t} + \theta^b_{B,t} = 0$, $c_{A,t} + c_{B,t} = D_t$, for all $t = 0, \ldots, T$ and (iii) the initial conditions $\theta^s_{A,-1} + \theta^s_{B,-1} = 1$ and $\theta^b_{A,-1} + \theta^b_{B,-1} = 0$ are satisfied.

The Euler equations that determine the equilibrium prices p^s_t and p^b_t of the stock and the bond, for $i = A, B$ and $t = 0, \ldots, T-1$, are:

$$p^s_t = \mathbb{E}_{i,t}\left[\beta \frac{u'_i(c_{i,t+1})}{u'_i(c_{i,t})}\left(p^s_{t+1} + D_{t+1}\right)\right] \tag{28}$$

$$p^b_t = \mathbb{E}_{i,t}\left[\beta \frac{u'_i(c_{i,t+1})}{u'_i(c_{i,t})} 1\right], \tag{29}$$

with $p^s_T = 0$. The one period riskless rate at time t is $r_t = 1/p^b_t - 1$.

Parameter uncertainty and learning implies that agents' posterior beliefs are state variables in the model, in addition to the level of the observed dividend. Furthermore, because of agents' heterogeneity, the distribution of consumption becomes a key state variable for the construction of an equilibrium. In sum, at each time, the equilibrium is characterized by (i) the dividend realization D_t, (ii) the dividend mean estimate, m_t, (iii) the dividend variance estimate, v^2_t, (iv) agents' consumption share, $\omega_{A,t} \equiv c_{A,t}/D_t = 1 - \omega_{B,t}$, which uniquely determines agents' individual $t+1$ state prices.

3.5 Solution methodology

Because the model does not admit an analytical solution, we rely on a numerical procedure to construct and analyze the equilibrium. In each decision node, the system of conditions (25)–(29) determines the equilibrium asset prices and agents' consumption. The intertemporal Euler's conditions (28)–(29) make the system simultaneously "backward" and "forward" at each point in time. The typical solution technique for this type of problems relies on forecasting functions and forward-backward iterations, e.g., Krusell and Smith (1998). Such an approach, however, does not extend naturally to a model of learning. Instead, we rely on the "time shift" recursive approach of Dumas and Lyasoff (2012). This approach simultaneously solves for (i) time t portfolios and prices and (ii) time $t+1$ consumption. As a result, the problem can be solved entirely through backward iteration.

To implement this approach, we shift the budget constraint (25) forward in time which, at time t, allows us to solve for agents' time-$t+1$ consumption and time-t portfolios. Unlike Dumas and Lyasoff (2012), who model uncertainty as a binomial tree, we characterize the state space by a discrete multidimensional grid consisting of four state variables $(D_t, m_t, v_t^2, \omega_{A,t})$ plus time. Because at each point in time there are fewer non-redundant securities than future possible states of the world, markets are incomplete. Therefore, agents are forced to trade in response to endowment shocks and/or updated beliefs. The core of the algorithm consists in recursively constructing time $t+1$ individual investors' consumption (state prices) as functions of time t's consumption. Appendix C provides a detailed description of our solution algorithm.

3.6 Results

In this section, we illustrate the properties of the equilibrium portfolio flows and asset prices emerging from our dynamic heterogeneous agent model. The main purpose of this analysis is to illustrate that the intuition of the simple model of Section 2 survives in a full-fledged general equilibrium model. In addition, however, the model of this section allows us to investigate the effect of agents' heterogeneity on the equilibrium risk-free rate and the consumption process, an analysis that is not possible in the simple model of Section 2. Finally, the multi-period setting of this section allows to analyze optimal impulse response functions of consumption and investment following dividend surprises.

Parameters. We assume an economy where dividends are IID normal random variables with mean $\mu = 1$ and volatility $\sigma = 0.1$, $D_t \sim \mathcal{N}(1, 0.1^2)$, and agent have a $T = 5$ years planning horizon. At $t = -1$, both agents are endowed with half of the risky asset and no bonds. Both agents start with the correct priors for the mean and variance of D (i.e., $m_{-1} = 1$ and $v_{-1}^2 = 0.01$), and both have the same CARA utility with a risk-aversion of $\gamma_A = \gamma_B = 1$.[25] We assume a time preference parameter $\beta = 0.9$, a constant gain parameter $a = 0.05$ and a coefficient of ambiguity aversion $\kappa = 0.15$.[26] To better understand the effect of learning in equilibrium, we analyze the equilibrium response of agents' portfolio holdings and asset prices associated to surprises in the dividend realizations. Specifically, we follow the equilibrium quantities of interest along a dividend path that starts at a level equal to the

[25]Because both agents own half of the tree and the price of the risky asset in the different settings is approximately 3.67, the agents' implied coefficient of relative risk aversion is $\gamma^{\text{RRA}} = \frac{1}{2} \times 3.67 \approx 1.8$.

[26]The choice of a moderate level of the ambiguity aversion parameter κ allows us to focus on a setting where both types of agents participate actively in the market, as in the upper panels of Figure 2.

true mean $d_0 = 1$, exhibits a temporary shock at time $t = 1$, and returns to its true mean, $D_t = 1$, for the remaining periods.[27]

[**Figure 3 about here.**]

Portfolios. Figure 3 shows the optimal stock (left panel, θ_B^s) and bond (right panel, θ_B^b) holdings of the Bayesian investor B as a function of the dividend surprise e at $t = 1$. The ambiguity-averse investor A holds $\theta_A^s = 1 - \theta_B^s$ units of the risky asset and $\theta_A^b = -\theta_B^b$ units of the risk-free asset. As expected, the investment in the risky asset of agent A declines significantly in ambiguity aversion κ. As in the simplified model of Section 2, Bayesian investors *reduce* their holdings in the risky asset after sufficiently large dividend surprises for all values of the ambiguity aversion parameter $\kappa > 0$. The magnitude of this effect increases in the level of ambiguity aversion. For $\kappa = 0$, dividend surprises do not induce any flows. The right panel of Figure 3 shows that in equilibrium Bayesian agents borrow from the ambiguity averse, $\theta_B^b < 0$, and that the borrowing position increases with the level of ambiguity aversion κ. As in the simple model of Section 2, in the absence of ambiguity aversion, $\kappa = 0$, there are no portfolio flows in equilibrium.

Risk premium and risk-free rate. Figure 4 reports the equilibrium risk premium (left panel) and risk-free rate (right panel) as a function of the dividend surprise e at $t = 1$. As for the case of the simple model of Section 2, the risk premium increases with both positive and negative dividend surprises and is larger for higher values of the ambiguity aversion parameter κ. In contrast, the equilibrium risk-free rate is declining in the size of the dividend surprise and largely unaffected by the presence of ambiguity aversion. A large dividend surprise, makes agents better off. This reduces the need to borrow to finance holdings in the risky asset and lowers the risk-free rate. Intuitively, a large dividend surprise increases the desire to transfer consumption to future states, hence, inflates asset prices and reduces the interest rate.

[**Figure 4 about here.**]

Consumption. Figure 5 reports the equilibrium consumption path of the Bayesian (left panel) and ambiguity averse (right panel) agents, following a positive (solid lines) or negative

[27]Note that, in general, for $t > 1$, $D_t = 1 \neq m_t$, since agents learn from the dividend shock at $t = 1$.

(dashed line) dividend surprise at $t = 1$. The colored lines correspond to different magnitudes of the shock. The black solid line represents the equilibrium consumption path in the absence of surprises. We set the coefficient of ambiguity aversion to $\kappa = 0.15$. Without dividend surprises, the black lines indicate that the equilibrium consumption of the Bayesian investor slightly increases over time, while that of the ambiguity averse investor slightly declines. Because the Bayesian agent holds a larger position in the risky asset, his consumption is more sensitive to the dividend surprise at $t = 1$ than that of agent A.

Interestingly, while the consumption of the B agent in the positive surprise paths remains above the consumption level of the no-surprise path, the consumption of the ambiguity averse agent exhibits *over-reaction* and stabilizes at a level *below* the no-surprise path. This finding indicates that type-A agents benefit from negative dividend surprises. To understand this effect, recall that agents earn the *objective* risk premium $\lambda^{\text{obj}} = \mu - m(e) + \lambda^{\text{subj}}(e)$, given in equation (17), but act on the basis of their *subjective* (observable) risk premium $\lambda^{\text{subj}}(e)$. Negative surprises imply an estimated mean $m(e)$ that is below the true value μ. Therefore, $\lambda^{\text{obj}} > \lambda^{\text{subj}}(e)$. Since ambiguity averse agents trade relatively more aggressively following surprises, after a negative dividend surprise their equilibrium consumption settles at a level higher than the no-shock consumption (black solid line). The effect is symmetric for positive dividend surprises.

[Figure 5 about here.]

To understand the consumption dynamics of Figure 5, in Figure 6 we report the equilibrium impulse response function for the risky asset holdings of the Bayesian (left panel) and the ambiguity averse (right panel) agents. Both panels show the equilibrium portfolio *deviations* from the holdings in the absence of dividend surprises. Consistent with the intuition of the stylized model of Section 2, ambiguity averse agents *increase* their holdings of the risky asset following negative surprise (dashed lines) and sufficiently large positive surprises (red solid line) while Bayesian agents decrease their holdings. Higher portfolio holdings imply a higher sensitivity of consumption to dividend surprises and generate the over-reaction of equilibrium consumption of ambiguity averse agents in the right panel of Figure 5.

Following a bad signal at time $t = 1$, Bayesian agents reduce their holdings in the risky asset (left panel, dashed red line). A negative signal implies a downward revision of the estimated mean $m_t(e)$, an upward revision in the estimated volatility $\hat{\sigma}_t$ and an increase in both the objective and subjective risk premium (see, equation (17)). Because in the subsequent periods, the dividend realizations are set at their true mean, $D_t = \mu$ for $t > 2$, they represent relatively good news for Bayesian investors who gradually reduce their short positions relative to the no-surprise path, as illustrated by the dashed lines. In contrast, following a good signal at time $t = 1$, the subjective risk premium increases because of an upward revision in the volatility $\hat{\sigma}_t$ but the objective risk premium *decreases*, because investors revise upward the mean estimate $m_t(e)$. In the periods following the $t = 1$ shock, the dividend realizations are set at their true mean, $D_t = \mu$ and therefore represent relatively bad news for Bayesian investors who gradually reduce their long positions, as illustrated by the solid lines in the left panel.

[**Figure 6 about here.**]

4 Empirical analysis

In this section, we provide evidence in support of our model predictions. Two challenges arise when bringing the model to the data: (i) how to map the idealized agent types in our model to observable classes of market participants. (ii) how to find good empirical characterizations of surprising changes in future dividend prospects.

With regard to the first challenge, when interpreting the empirical results, we take the ambiguity averse, type-A, agents of our model as representatives of the class of individual investors and type-B agents as representative of the class of institutional investors. This classification is admittedly crude, given the substantial heterogeneity observed within each investor type. It is however motivated by a large body of empirical and experimental evidence that favors the interpretation of individual investors being relatively more averse to uncertainty than institutions. For instance, Li, Tiwari, and Tong (2017) provide empirical support for the assumption that retail investors have a stronger desire for robustness. Moreover, experimental studies document that ambiguity aversion is influenced by the perceived competence of decision makers (known as *competence hypothesis*, see Heath and Tversky, 1991), or "by a comparison with less ambiguous events or with more knowledgeable individuals" (known as *comparison hypothesis*, see Fox and Tversky, 1995). Relatedly, Graham, Harvey, and Huang (2009) argue that investors who perceive themselves competent are likely to have less parameter uncertainty about their subjective distribution of future asset returns. Because institutional investors have typically access to larger resources and are professional investors, they might therefore be perceived by individuals as more knowledgeable.

With regard to the second challenge, we use exceptionally high or low market returns as a timely signal on which agents condition their expectations about future dividend payments. In the model, agents use dividend payments as signals of future expected profitability. Ideally, unexpected firms' earnings would be a natural measure of changes in profitability. However, earnings reports are notoriously noisy and contain outdated information. The use of returns as indicators for news about profitability is justified by our model, in which realized dividends and contemporaneous price reactions are highly correlated in equilibrium.

Within this framework, we provide empirical evidence of the two main predictions of our model: (i) Exceptionally good or bad news about future corporate profitability lead to an increase in corporate ownership by individual investors and a corresponding decrease of holdings by institutional investors; (ii) Using only in-sample data, investors' estimate of the expected risk premium around surprising signals about corporate profitability are higher than on average.

We conduct our analysis out-of-sample, that is, from the perspective of investors who learn with fading memory as they observe dividend realizations over time. Specifically, we are interested in estimating an empirical counterpart of the premium $\lambda(e)$ in equation (17) which is observable by investors in real-time.

4.1 Data

We use two different data sources: (i) aggregate level and flow data on corporate equity holdings of households and the domestic financial sector from 1952.Q1 to 2020.Q4, obtained from the Federal Reserve of St. Louis database (FRED)[28] and (ii) institutional holdings of U.S. firms from 2000.Q1 to 2020.Q1, obtained from the Thomson Reuters OP Global Ownership database (Consolidated Holdings), which we augment with information from Compustat-CapitalIQ. From CRSP we obtain return data of all firms listed at NYSE, AMEX, and NASDAQ from 1965.01 to 2020.12. To measure surprises in firms' future profitability we rely on standarized market returns, z_r, i.e., normalized to have zero mean and unit variance using a rolling window of 20 quarters. The market return is taken from Kenneth French's data library.[29]

[28] Data source: `https://fred.stlouisfed.org/tags/series`.
[29] Data source: `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

We use level and flow data of corporate equities held by households and by components of the domestic financial sector, according to the FRED definitions: mutual funds, security brokers and dealers, closed-end and exchange-traded funds, other financial business, private depository institutions, insurance companies and pension funds, and monetary authority. Given the inertia in pension funds portfolio allocation, (see e.g. Agnew, Balduzzi, and Sunden, 2003; Hu, McLean, Pontiff, and Wang, 2014), we do not consider pension fund data in our analysis. From the level and flow data of households and the financial sector, we compute quarterly aggregated (value-weighted) equity returns.

The Thomson Reuters OP Global Ownership (Consolidated Holdings) database covers 13F reporting institutions, mutual, pension and insurance funds, declarable stakeholders and UK share registers. After excluding firms with market cap below $5 millions, we end up with quarterly data for the time span 2000.Q1-2020.Q1 for 8,488 firms with 274,697 firm-quarter observations.[30]

For the subsequent analysis, we use standardized quarterly market returns, z_r, i.e., normalized to have zero mean and unit variance using a 20-quarter rolling windows to group observations in both data sets into five bins. The breakpoints for bins are given by the 7.5%, 25%, 75%, and 92.5% percentiles of a standard normal distribution.

4.2 Equilibrium flows

Figure 7 shows the relationship between changes in institutional ownership, $\Delta\theta_B^s$, and standardized market returns, z_r. The left panel plots aggregate data from the FRED database, the right panel shows changes in ownership of individual firms from the Thomson Reuters OP Global Ownership database for firms listed on the NYSE, AMEX or NASDAQ exchages and with market capitalizations in excess to $5 million. Mean values are in black, median values in red. As the figure shows, exceptionally bad as well as exceptionally good returns, representing signals of extraordinary negative and positive news about corporate profitability, are associated with low or even negative changes in institutional ownership. In contrast, neutral signal realizations, indicating lack of surprise, exhibit an increase in institutional

[30]As standard in this strand of literature, outstanding shares not held by institutional investors are assumed to be held by private investors.

ownership. These changes should be interpreted relative to the substantial trend towards institutional ownership which is present since 1980, see Stambaugh (2014).[31]

[**Figure 7 about here.**]

While the reduction of institutional ownership in response to negative surprises is in line with the ample evidence about private investors acting as liquidity providers who meet institutional investors' demand for immediacy (see, e.g., Kaniel, Saar, and Titman, 2008; Barrot, Kaniel, and Sraer, 2016; Glossner, Matos, Ramelli, and Wagner, 2020; Pástor and Vorsatz, 2020), this line of reasoning would not explain the reduction in θ_B^s after positive surprises found in both data sets.

[**Table 1 about here.**]

Table 1 provides details on the analysis underlying the results in Figure 7. The table shows that high as well as low standardized returns z_r are associated with low contemporaneous changes in institutional ownership, $\Delta\theta_B^s$, intermediate z_r-values come with an increase in institutional ownership $\Delta\theta_B^s$. The results hold regardless of whether we consider the mean or the median changes within bins. In the upper panel we show results for the FRED data set. In order to corroborate our claim that institutional ownership declines after good and bad surprises, we conduct a non-parametric Kruskal-Wallis (KW) rank sum test (Kruskal and Wallis, 1952).[32] The KW test confirms that $\Delta\theta_B^s$ differs across bins, and the post-hoc Dunn test attests that central bins have a significantly higher $\Delta\theta_B^s$ compared to the extreme bins (see the corresponding p-values).

The lower panel shows results for the Thomson Reuters Global Corporte Ownership data. Since individual firm observations are correlated within each quarter, we perform a

[31] In the FRED data, institutional ownership (excluding pension funds) increases from 3% in 1952.Q1 to 42% in 2020.Q4. In the individual-firm data, over the sample period from 1999.Q1 to 2020.Q1 institutional ownership increases from 32% to 59% for firms with market capitalization in excess to \$5 millions, and from 44% to 76% for firms with market capitalization exceeding \$1 billion. Hence, quarterly changes in institutional ownership must be compared to the average growth of institutional ownership (approximately 14bp per quarter for FRED data, and 30bp per quarter for our individual-firm data).

[32] The KW test, an extension of the (Wilcoxon)-Mann-Whitney U-test, is a non-parametric rank-sum test analyzing whether observations in the different bins originate from the same distribution. While the test indicates whether observations in one bin are different from observations in the others bins, it does not indicate which bins cause this results. For that purpose, a subsequent (post hoc) Dunn test (Dunn, 1964) allows for a pairwise comparison of the bins.

clustered Wilcoxon rank sum test (clustered by quarter) to conduct the pair-wise comparison between bins. The results confirm the findings in the FRED data set. Change in institutional ownership in the first bin (low market returns) is significantly lower than in bins 3 and 4. Change in institutional ownership in bin 5 is significantly lower than in bin 4.

4.3 Equilibrium risk premium

Our second model prediction is the U-shaped relationship between news and risk premia. The left panel of Figure 8 shows estimates of the equity risk premium, computed as return in excess of the 3-month T-Bill rate, from aggregate FRED data. The right panel shows estimates of the market risk premium from a conditional Fama-MacBeth regression using return data of all stocks (common equity) traded on NYSE, AMEX or NASDAQ from CRSP in the period from 1965 to 2021. Specifically, we first compute asset βs through time series regressions of individual monthly returns in excess to the 1-month T-Bill rate on the value-weighted market excess return over a sliding window of 36 months. We then estimate cross-sectional regressions of individual quarterly excess returns on these β-estimates (see Fama and MacBeth, 1973). The slope coefficients of these regressions, i.e., the quarterly estimates of the market risk premium, are then sorted into bins conditional on the lagged market return. Hence, the mean and the median coefficient within each bin represent estimates of the expected compensation per unit of market risk exposure conditional on lagged returns.

Consistent with Cao, Wu, and Wu (2022), we find that the low-beta anomaly, that is, the negative relationship between equity beta and the risk premium, is present during times of low uncertainty when the standardized market returns z_r are close to zero. In contrast, during times of high uncertainty, i.e., z_r far away from zero, there is a positive premium for bearing market risk, implying that a "betting-against-beta" strategy would not be profitable. The premium reported in the right panel of Figure 8 does not include the cross-sectional regression intercept, hence, it should be interpreted as an estimate of the *marginal* premium offered for bearing one additional unit of market β risk rather than the total expected premium for holding the market portfolio. While the expected marginal premium is even negative in calm times, consistent with the low-beta anomaly, the total premium for holding the market is positive, since the intercept is significantly positive under these conditions (a fact also reported by Cao, Wu, and Wu, 2022).

[**Figure 8 about here.**]

Both panels of Figure 8 show that market risk premia are higher following negative *as well as* positive surprises. This finding is partly in line with Nagel (2012), who shows that the returns of a short-term reversal strategy can be interpreted as a proxy for the returns of liquidity provision, and that these returns are especially high during periods of financial turmoil. However, while this argument can rationalize the higher returns in the left tail of the distribution of return surprises, it is silent about the observed high premia following positive return surprises.

[**Table 2 about here.**]

Table 2 provides details on the analysis underlying the results in Figure 8. The table shows that while high as well as low lagged standardized returns z_r are associated with high risk premia, intermediate lagged z_r-values imply low risk premia. The results hold for the mean as well as for the median premium within bins. The upper panel shows results for FRED data while the lower panel shows results for CRSP data. In order to test the claim that risk premia increase after good and bad surprises, in the right side of both panels we report a non-parametric Kruskal-Wallis rank sum test. The tests confirm that risk premia differ across bins, and subsequent post-hoc Dunn tests show that central bins have significantly lower risk premia compared to the extreme bins, as indicated by the corresponding p-values.

5 Conclusion

We explain asset prices and portfolio flows following episodes of increased economic uncertainty using an equilibrium model in which agents learn about the mean and the volatility of the endowment process and differ in their concerns about parameter uncertainty. We show that, in equilibrium, ambiguity averse investors hold more conservative portfolios but trade more aggressively in response to surprises about corporate profitability. Regardless of the sign of the surprise—positive or negative—Bayesian investors reduce their share in the risky asset while ambiguity averse investors increase their share. Agents' learning about volatility gives rise to a time-varying equilibrium risk premium. While in equilibrium innovations to the expected dividend are immediately absorbed in prices, large positive and negative surprises generate upward revisions in the estimated dividend volatility and increase risk premia. Because ambiguity averse preferences exhibit first-order risk aversion, the equilibrium risk premium depends linearly on both the variance and the volatility of the endowment. When the estimated volatility increases, as it happens following dividend surprises, the linearity in volatility makes the risky asset relatively more attractive to ambiguity averse agents who increase their risky holdings, compared to ambiguity-neutral agents.

We first illustrate these results in a simple one-period model which is analytically tractable. We then analyze a multi-period general equilibrium endowment economy with intertemporal consumption. When agents learn with fading memory about the mean and the variance of the endowment process, uncertainty is time-varying and persists over time. The multi-period model allows us to generalize the intuition of the simple model to a larger setting and to analyze the dynamics of optimal consumption and portfolio responses to dividend surprises. To solve this general model, we extend the incomplete-market backward approach of Dumas and Lyasoff (2012) to account for learning and heterogeneity in the preference for robustness.

Three main ingredients are needed to explain flows of funds and risk premia in our setting: (i) differences in the preference for robustness; (ii) learning about variance; and (iii) market clearing. Without ambiguity aversion, agents do not rebalance their portfolio after surprises but only risk premia react. Without learning about dividend variance, the equilibrium risk premium is a constant and agent's portfolios are static. Similar to prior work in the literature on learning and predictability, risk premia in our model are counter-cyclical. However, in contrast to studies that assume a known variance, in our setting a part of the risk-premium is observable to forward-looking investors. From an econometrician's perspective, this implies that good and bad surprises have an asymmetric effect on the objective risk premium.

Finally, we bring the predictions of our model to the data by analyzing portfolio holdings of institutional and individual investors. Using aggregated data from FRED as well as single stock data from the CRSP-Compustat universe, we provide evidence that institutional investors tend to reduce their share in corporate ownership when indicators of future corporate profitability are exceptionally bad and exceptionally good. We further find that the expected risk premium is higher after both positive and negative surprises. These findings are consistent with the predictions of our model when institutions trade with ambiguity averse investors who are conservative in their holdings but aggressive in their trades.

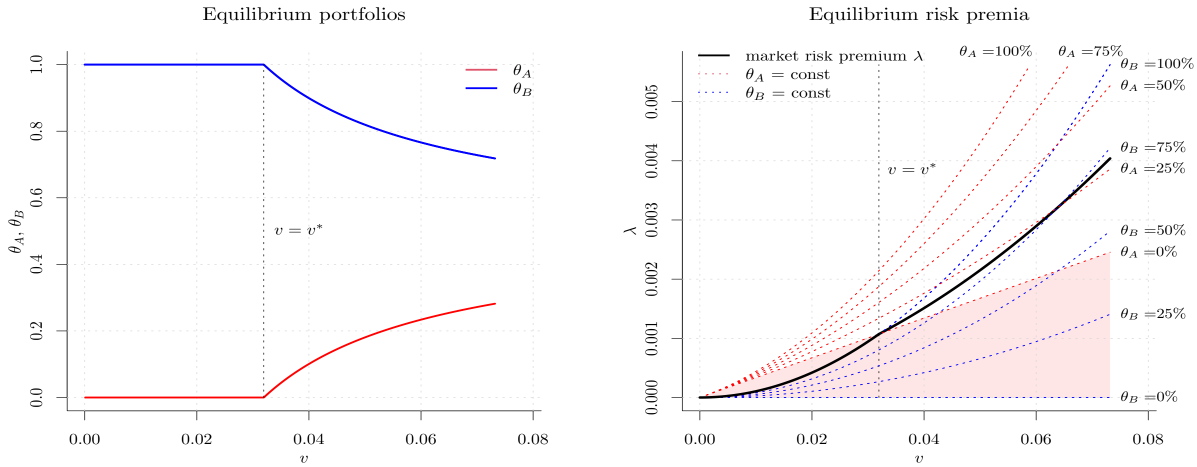

Figure 1: **Equilibrium portfolios and risk premia.** The left panel shows the equilibrium portfolios θ_A and θ_B as a function of the dividend volatility v. The right panel shows iso-θ lines of agent A (dotted red) and B (dotted blue). These lines represent the set of values (v, λ) for which the equilibrium portfolios in equations (3) and (6) corresponds to a given fractional holding of the endowment tree. The locus of points (v, λ) at which market clears, $\theta_A + \theta_B = 1$ identify the equilibrium risk premium λ as a function of the dividend volatility v (solid black line). The vertical dashed line indicates the participation threshold $v^* \equiv \frac{\sqrt{n}}{n+1} \frac{\kappa}{\gamma_B}$. Parameter values: $n = 20$, $\gamma_A = \gamma_B = 1$, $\kappa = 0.15$.

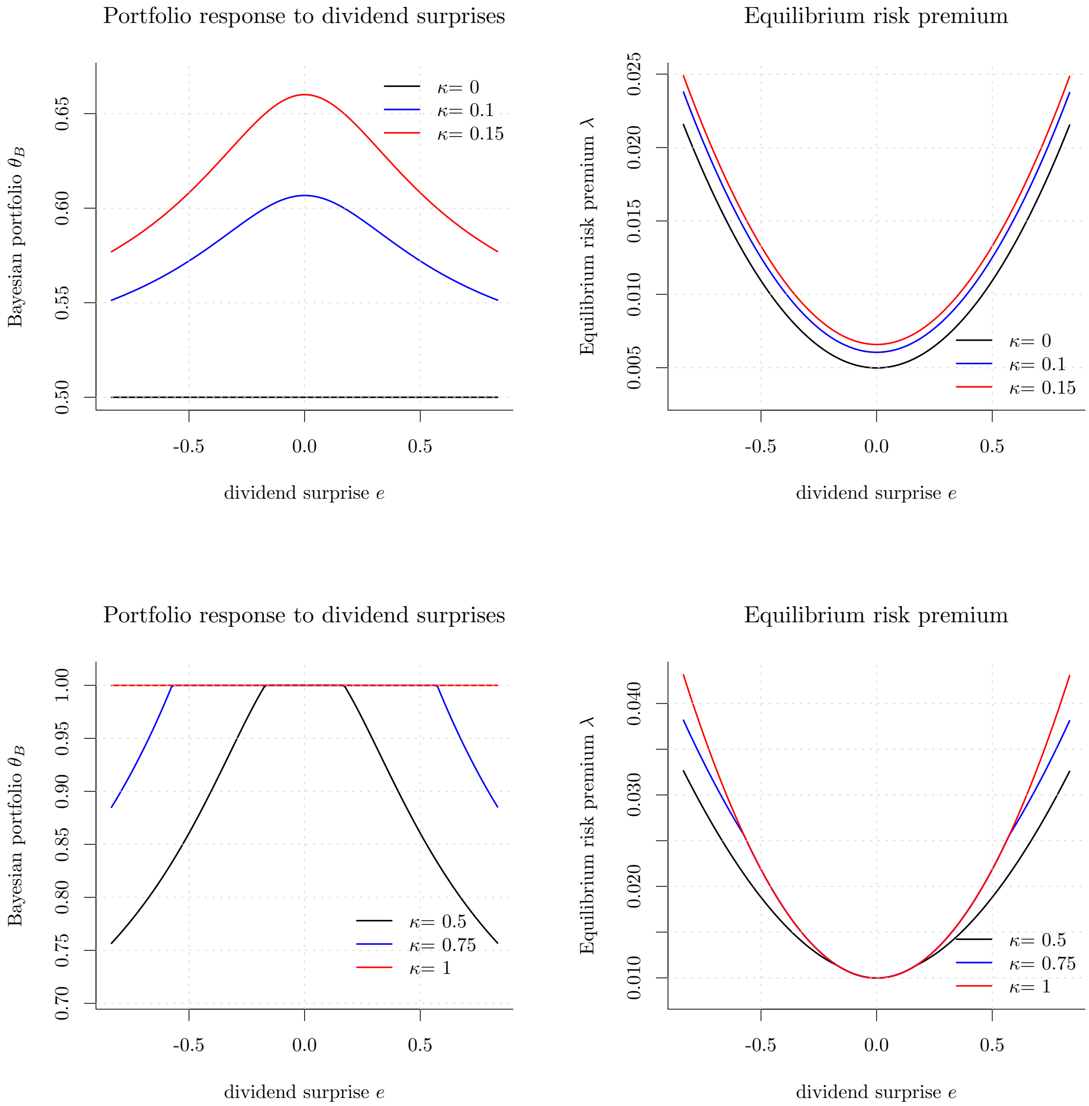

Figure 2: **Equilibrium portfolios and risk premia.** The figure shows optimal holdings of the risky asset for the Bayesian investor θ_B (left panels) and equilibrium subjective risk premia λ (right panels) as a function of dividend surprises $e = D - m$. The top panels consider low levels of the ambiguity aversion parameter κ while the bottom panels considers high levels of κ. Parameter values: $n = 20$, $\gamma_A = \gamma_B = 1$, $m = 1$, $v^2 = 0.01$.

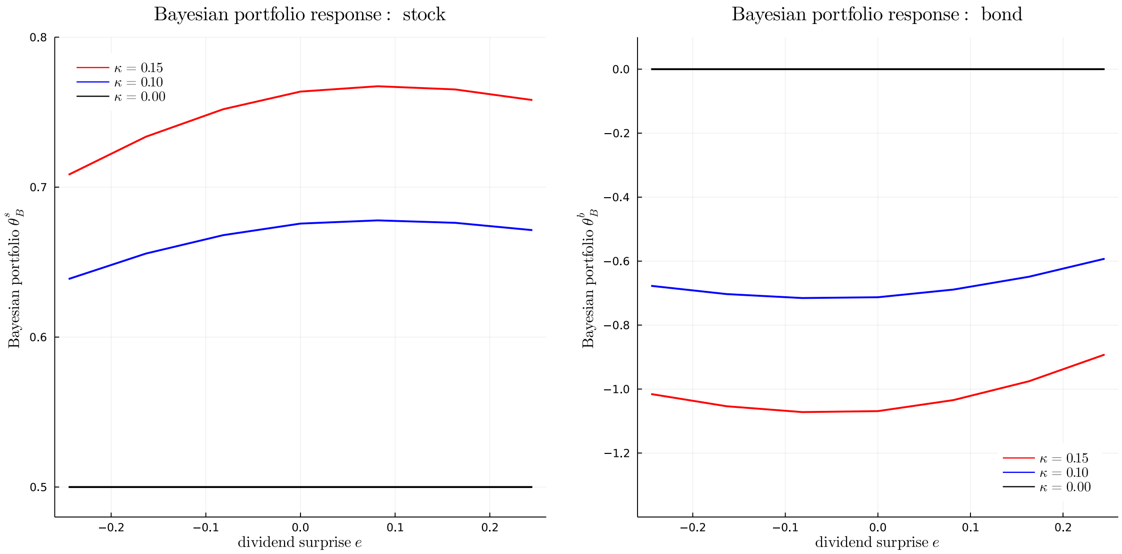

Figure 3: **Equilibrium stock and bond portfolios: dynamic model.** The figures reports the Bayesian portfolio holding of stocks (left panel) and bonds (right panel) as a function of the dividend surprises $e = D - m$ at $t = 1$. Parameter values: $T = 5$, $a = 1/20$, $\gamma_A = \gamma_B = 1$, $\beta = 0.9$, $m_0 = 1$, $v_0^2 = 0.01$, $\theta^s_{A,-1} = \theta^s_{B,-1} = 0.5$, $\theta^b_{A,-1} = \theta^b_{B,-1} = 0$.

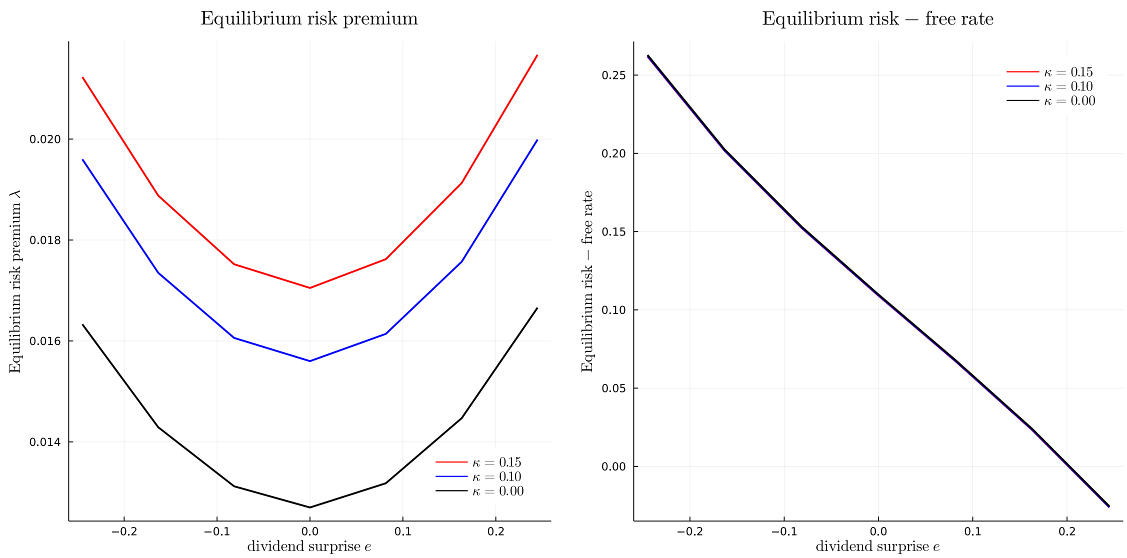

Figure 4: **Equilibrium risk premium and risk-free rate: dynamic model.** The figure reports the equilibrium risk premium (left panel) and risk-free rate (right panel) as a function of the dividend surprise $e = D - m$ at $t = 1$. Parameter values: $T = 5$, $a = 1/20$, $\gamma_A = \gamma_B = 1$, $\beta = 0.9$, $m_0 = 1$, $v_0^2 = 0.01$, $\theta^s_{A,-1} = \theta^s_{B,-1} = 0.5$, $\theta^b_{A,-1} = \theta^b_{B,-1} = 0$.

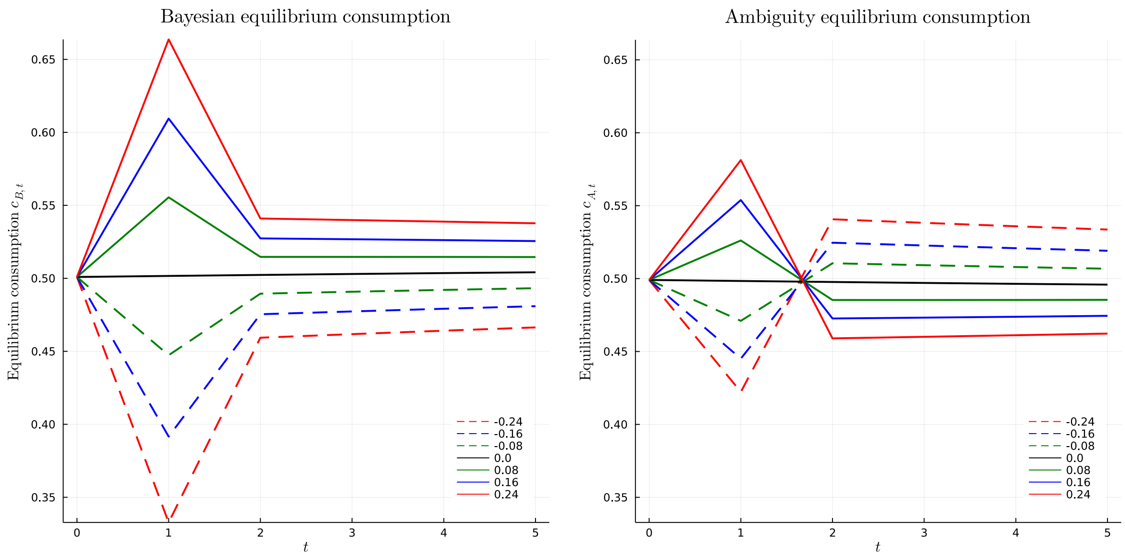

Figure 5: **Equilibrium consumption dynamics.** The figure reports the impulse response function of equilibrium consumption of the Bayesian (left panel) and ambiguity averse (right panel) agents, following a positive (solid lines) or negative (dashed line) surprise shock $e_1 = 0, \pm 0.08, \pm 0.16, \pm 0.24$. The colored lines corresponds to different magnitudes of the shock. The black solid line represent the equilibrium consumption path in the absence of surprises. The ambiguity parameter is set to $\kappa = 0.15$. Parameter values: $T = 5$, $a = 1/20$, $\gamma_A = \gamma_B = 1$, $\beta = 0.9$, $m_0 = 1$, $v_0^2 = 0.01$, $\theta^s_{A,-1} = \theta^s_{B,-1} = 0.5$, $\theta^b_{A,-1} = \theta^b_{B,-1} = 0$, $D \sim \mathcal{N}(1, 0.1^2)$.

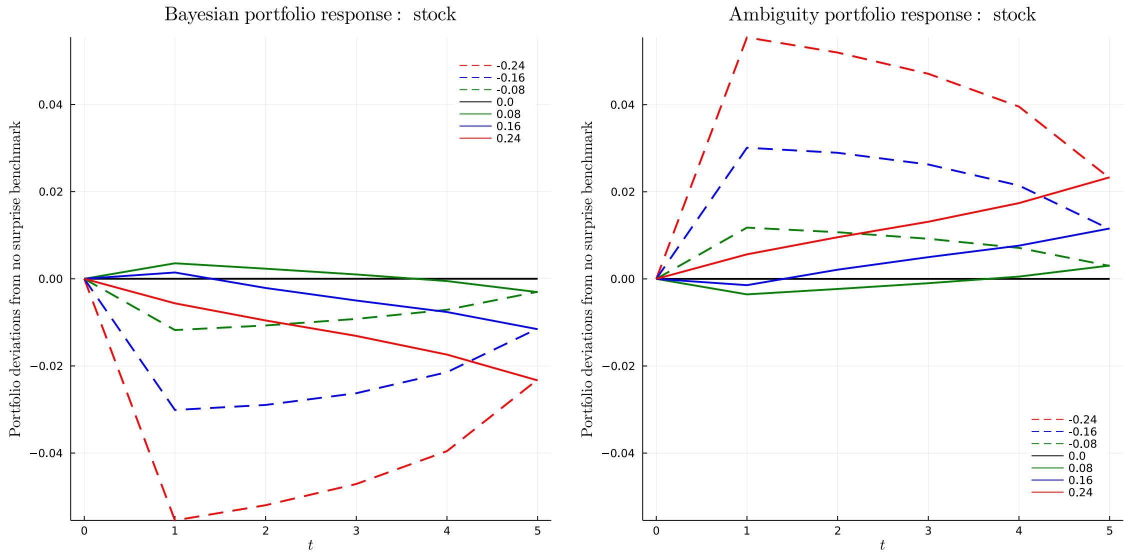

Figure 6: **Equilibrium portfolio dynamics.** The figure reports the impulse response function of equilibrium risky asset holdings of the Bayesian (left panel) and ambiguity averse (right panel) agents, following a positive (solid lines) or negative (dashed line) surprise shock $e_1 = 0, \pm 0.08, \pm 0.16, \pm 0.24$. The colored lines correspond to different magnitudes of the shock and represent portfolio *deviations* from the portfolio holding in the absence of surprises. The ambiguity parameter is set to $\kappa = 0.15$. Parameter values: $T = 5$, $a = 1/20$, $\gamma_A = \gamma_B = 1$, $\beta = 0.9$, $m_0 = 1$, $v_0^2 = 0.01$, $\theta^s_{A,-1} = \theta^s_{B,-1} = 0.5$, $\theta^b_{A,-1} = \theta^b_{B,-1} = 0$, $\tilde{D} \sim \mathcal{N}(1, 0.1^2)$.

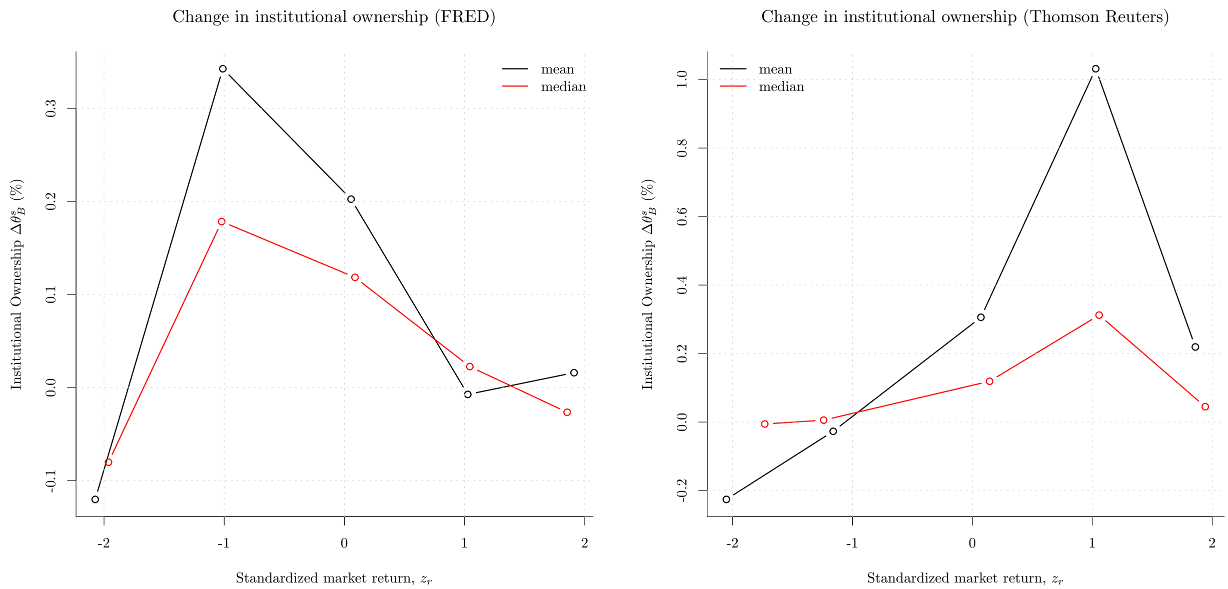

Figure 7: **Change in institutional holdings and dividend surprises.** The figure shows mean (black) and median (red) quarterly changes in institutional ownership, $\Delta\theta_B^s$, as a function of dividend surprises. As a proxy for surprises we use the standardized quarterly market returns, z_r, obtained from a 20-quarter rolling window. We use z_r to group observations into five bins with breakpoints given by the 7.5%, 25%, 75%, and 92.5% percentiles of a standard normal distribution. The left panel shows results for data from the Federal Reserve Bank of St. Louis database. Ownership data are calculated from equity level data of market participants (households and financials). The right panel shows results for individual firms listed on NYSE, AMEX or NASDAQ. Ownership data are from Thomson Reuters Global Ownership database restricted to common shares traded on NYSE, AMAX or NASDAQ with market capitalization larger than $5 millions. The market return is taken from Kenneth French's data library.

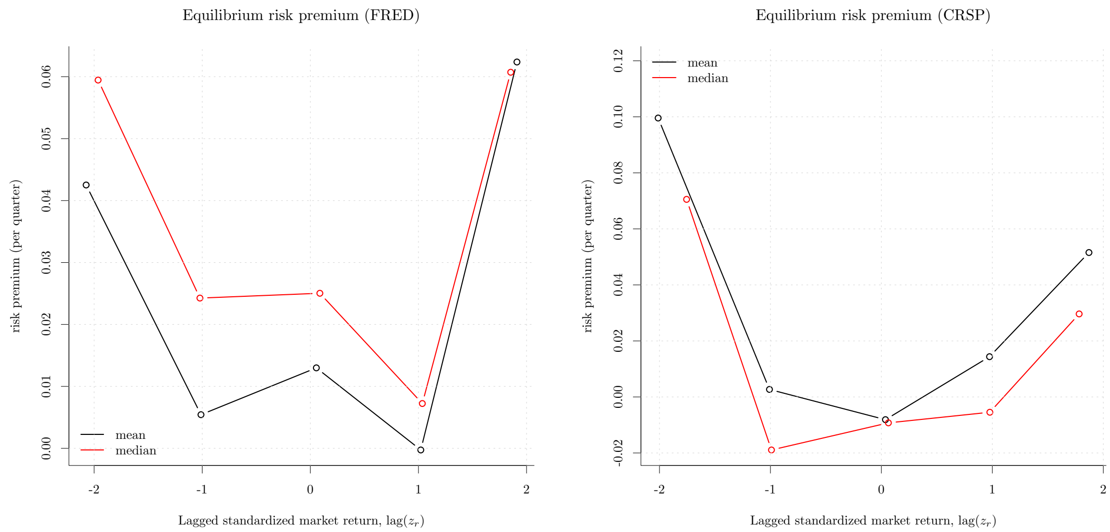

Figure 8: **Risk premia and dividend surprises.** The figure shows the mean (black) and median (red) market risk premia as a function of dividend surprises. As a proxy for surprises, we use the standardized quarterly market returns, z_r, obtained from a 20-quarter rolling window. We use z_r to group observations into five bins with breakpoints given by the 7.5%, 25%, 75%, and 92.5% percentiles of a standard normal distribution. The left panel shows results for data from the Federal Reserve Bank of St. Louis database. Return data are calculated from equity level and flow data of market participants (households and financials), and the risk premium is computed as excess return over the 3-month T-Bill rate. The right panel shows the conditional beta premium calculated from Fama-MacBeth regressions of returns of common shares traded on NYSE, AMAX or NASDAQ with market capitalization larger than $5 millions. The market return is taken from Kenneth French's data library.

FRED										
	mean		median			KW: $\chi^2 = 13.93$, $p = 0.0075$				
bin	z_r	$\Delta\theta^s_B$ (%)	z_r	$\Delta\theta^s_B$ (%)	n	Dunn post hoc				
						bin	1	2	3	4
1	-2.07	-0.12	-1.96	-0.08	24					
2	-1.01	0.34	-1.02	0.18	33	2	0.00			
3	0.06	0.20	0.09	0.12	146	3	0.01	0.24		
4	1.03	-0.01	1.04	0.02	36	4	0.36	0.02	0.06	
5	1.91	0.02	1.85	-0.03	18	5	0.42	0.05	0.18	0.98

Thomson Reuters Global Ownership										
	mean		median							
bin	z_r	$\Delta\theta^s_B$ (%)	z_r	$\Delta\theta^s_B$ (%)	n	clustered Wilcoxon rank sum				
						bin	1	2	3	4
1	-2.05	-0.23	-1.73	-0.01	37,457					
2	-1.16	-0.03	-1.24	0.01	23,774	2	0.51			
3	0.07	0.31	0.14	0.12	162,316	3	0.05	0.29		
4	1.03	1.03	1.06	0.31	40,832	4	0.01	0.05	0.05	
5	1.86	0.22	1.94	0.04	10,318	5	0.20	0.62	0.18	0.08

Table 1: **Change in institutional holdings and dividend surprises.** The table shows the relationship between institutional ownership $\Delta\theta^s_B$ (%) and dividend surprises z_r. As a proxy for surprises we use the standardized quarterly market returns, z_r, obtained from a 20-quarter rolling window. We use z_r to group observations into five bins with breakpoints given by the 7.5%, 25%, 75%, and 92.5% percentiles of a standard normal distribution. The number of observations in each bin is n, and $\Delta\theta^s_B$ (%) is the quarterly change in institutional ownership in percent. The top panel shows results for data from the Federal Reserve Bank of St. Louis database. Ownership data are calculated from equity level data of market participants (households and financials). Kruskal-Wallis (KW) tests for difference in median values of $\Delta\theta^s_B$ across the bins, and the post hoc Dunn test is used to conduct pairwise comparisons. The bottom panel shows results for individual firms with ownership data from Thomson Reuters Global Ownership database restricted to common shares traded on NYSE, AMAX or NASDAQ and market capitalization larger than $5 millions. The clustered Wilcoxon rank sum test clusters observations within the same quarter when performing bin-wise comparisons. The market return is taken from Kenneth French's data library.

FRED										
	mean		median			KW: $\chi^2 = 10.60$, $p = 0.0314$				
bin	lag(z_r)	$r - r_f$ (%)	lag(z_r)	$r - r_f$ (%)	n	Dunn post hoc				
1	-2.07	4.25	-1.96	5.95	24	bin	1	2	3	4
2	-1.01	0.54	-1.02	2.43	33	2	0.12			
3	0.06	1.30	0.09	2.50	146	3	0.08	0.87		
4	1.02	-0.03	1.04	0.72	35	4	0.03	0.47	0.27	
5	1.91	6.24	1.85	6.07	18	4	0.54	0.04	0.02	0.01

CRSP										
	mean		median			KW: $\chi^2 = 18.72, p = 0.0009$				
bin	lag(z_r)	$r - r_f$ (%)	lag(z_r)	$r - r_f$ (%)	n	Dunn post hoc				
1	-2.01	9.96	-1.76	7.05	22	bin	1	2	3	4
2	-1.01	0.27	-0.99	-1.89	27	2	0.00			
3	0.04	-0.81	0.06	-0.93	125	3	0.00	0.98		
4	0.98	1.44	0.98	-0.55	34	4	0.01	0.49	0.38	
5	1.87	5.16	1.78	2.96	16	5	0.49	0.04	0.02	0.12

Table 2: **Risk premia and dividend surprises.** The table shows the relationship between risk premia and dividend surprises, z_r. We use standardized quarterly market returns z_r (i.e. normalized to have zero mean and unit variance using a rolling window of 20 quarters) to group observations into five bins, and within each bin we calculate mean and median values. The breakpoints for bins are given by the 7.5%, 25%, 75%, and 92.5% percentiles of a standard normal distribution. The number of observations in each bin is n, and $r - r_f(\%)$ is the quarterly excess return in percent. Observations are according to lagged z_r into five bins, and within each bin we calculate mean and median values. The top panel shows results for data from the Federal Reserve Bank of St. Louis database. Return data are calculated from equity level and flow data of market participants (households and financials), and the risk premium is indicated as excess return over the 3-month T-Bill rate. Kruskal-Wallis (KW) tests for difference in the risk premia $r - r_f$ across the bins, and the post hoc Dunn test is used to conduct pairwise comparisons. The bottom panel shows conditional Fama-MacBeth estimates of the market risk premium obtained in the cross-section of firms listed on NYSE, AMEX or NASDAQ, with market capitalizations larger than $5 millions. The market return is taken from Kenneth French's data library.

A Slope of agents' iso-portfolio curves

We prove that the slope of the ambiguity-averse agents' iso-portfolio curves is flatter than those of the Bayesian agents. This implies that the marginal rate of substitution between risk premium and dividend volatility is lower for ambiguity-averse agents. From equation (6) we derive the risk premium that agents A require for holding a fraction θ_A of the risky asset (the iso-portfolio line) and its derivative with respect to dividend volatility v as

$$\bar{\lambda}_A = \frac{\kappa v}{\sqrt{n}} + \gamma_A \theta_A \left(\frac{n+1}{n}\right) v^2, \tag{A.1}$$

$$\frac{\partial \bar{\lambda}_A}{\partial v} = \frac{\kappa}{\sqrt{n}} + 2\gamma_A \theta_A \left(\frac{n+1}{n}\right) v. \tag{A.2}$$

From equation (3) we do the same for the iso-portfolio lines of agents B

$$\bar{\lambda}_B = \gamma_B \theta_B \left(\frac{n+1}{n}\right) v^2, \tag{A.3}$$

$$\frac{\partial \bar{\lambda}_B}{\partial v} = 2\gamma_B \theta_B \left(\frac{n+1}{n}\right) v > 0. \tag{A.4}$$

We prove that along the equilibrium risk premium λ in equation (8) the slope of $\bar{\lambda}_A$ is flatter than the slope of $\bar{\lambda}_B$. We need to prove

$$\frac{\partial \bar{\lambda}_A}{\partial v} < \frac{\partial \bar{\lambda}_B}{\partial v}. \tag{A.5}$$

Using equations (A.1)–(A.4), and $\theta_B = 1 - \theta_A$, this is equivalent to prove

$$\frac{\kappa}{\sqrt{n}} + 2\gamma_A \theta_A \left(\frac{n+1}{n}\right) v < 2\gamma_B (1 - \theta_A) \left(\frac{n+1}{n}\right) v, \tag{A.6}$$

or, rearranging,

$$2\left(\gamma_A + \gamma_B\right) \theta_A \left(\frac{n+1}{n}\right) v < 2\theta_B \left(\frac{n+1}{n}\right) v - \frac{\kappa}{\sqrt{n}}. \tag{A.7}$$

We restrict our analysis to the region where both agents are in the market, $v > \frac{\sqrt{n}}{n+1} \frac{\kappa}{\gamma_B}$, and substitute equilibrium portfolios weights from equation (9) into the above inequality. This yields

$$2\gamma_B\left(\frac{n+1}{n}\right)v - 2\left(\frac{\sqrt{n}\kappa}{(n+1)v}\right)\left(\frac{n+1}{n}\right)v \;<\; 2\gamma_B\left(\frac{n+1}{n}\right)v - \frac{\kappa}{\sqrt{n}}, \tag{A.8}$$

$$2\frac{\kappa}{\sqrt{n}} \;>\; \frac{\kappa}{\sqrt{n}}, \tag{A.9}$$

which is true for $\kappa > 0$ and $n < \infty$ independently of v.

B Bayesian updating with "Fading Memory"

We start from standard Bayesian filtering theory (e.g., West and Harrison, 2006) and use the fact that when learning the mean and variance of a standard normally distributed variable (the dividend) the t/inverse gamma family of distributions serves as a natural conjugate prior. In order to formalize the concept of fading memory, we employ a so-called *discount-factor* approach which assumes that the "surplus variance" which is imposed on the historical estimate is proportional to estimation uncertainty of the mean. Starting from a t/inverse gamma distributed prior at $t = 0$, the joint time $t-1$ posterior estimate of the expected dividend μ and its variance σ is again t/inverse gamma distributed

$$(\mu|t-1, D_{t-1}) \;\sim\; T_{n_{t-1}}\left[m_{t-1}, s_{t-1}^2\right], \tag{B.1}$$

$$\left(\frac{1}{\sigma^2}\middle| t-1, D_{t-1}\right) \;\sim\; \Gamma\left[\frac{n_{t-1}}{2}, \frac{n_{t-1}v_{t-1}^2}{2}\right], \tag{B.2}$$

with n_{t-1} denoting the number of degrees of freedom at time $t-1$. Fading memory affects the process through which the time-$t-1$ posterior is transformed into a time-t prior, that is,

$$(\mu|t, D_{t-1}) \;\sim\; T_{n_{t-1}}\left[m_{t-1}, \delta s_{t-1}^2\right],\ \delta > 1, \tag{B.3}$$

$$\left(\frac{1}{\sigma^2}\middle| t, D_{t-1}\right) \;\sim\; \Gamma\left[\frac{n_{t-1}}{2}, \frac{n_{t-1}v_{t-1}^2}{2}\right]. \tag{B.4}$$

The parameter $\delta > 1$ increases the variance of the prior on μ and hence lowers its information quality. It parsimoniously captures the extent of an agent's fading memory. A larger value of δ corresponds to a stronger extent of fading memory.

Bayesian updating of the beliefs upon observation of the dividend D_t follows the recurrence (see West and Harrison, 2006, Sections 4.6 and 6.3):

$$m_t = (1-a_t)m_{t-1} + a_t D_t, \tag{B.5}$$

$$s_t^2 = a_t v_t^2, \tag{B.6}$$

$$a_t = \frac{\delta s_{t-1}^2}{\delta s_{t-1}^2 + v_{t-1}^2} = \frac{\delta a_{t-1}}{\delta a_{t-1}+1}, \tag{B.7}$$

$$v_t^2 = v_{t-1}^2 + \frac{v_{t-1}^2}{n_t}\left[\frac{e^2}{\delta s_{t-1}^2 + v_{t-1}^2} - 1\right]$$

$$= \left(1 - \frac{1}{n_t}\right) v_{t-1}^2 + \frac{1}{n_t}\left(1-a_t\right) e^2, \tag{B.8}$$

$$e = D_t - m_{t-1}, \tag{B.9}$$

$$n_t = n_{t-1} + 1. \tag{B.10}$$

From equation (B.7) we note that the gain a_t follows a deterministic recurrence. With unrestricted memory, i.e., $\delta = 1$, a_t converges asymptotically to 0. Thus, in this case, agents' beliefs m_t, s_t^2 and v_t^2 become over time less and less responsive to new observations D_t, and agents learn the moments of the dividend process with perfect precision. In contrast, with fading memory, $\delta > 1$, the Bayesian gain a_t remains positive in the limit, and we have

$$\lim_{t\to\infty} a_t = a = \frac{\delta - 1}{\delta} =: \frac{1}{n_{\text{eff}}} > 0. \tag{B.11}$$

Asymptotically, the updating recurrence (B.5) for m_t stays responsive to new observations as if at any time t there are in total only n_{eff} dividend observations available. Hence, we can interpret this case as learning with fading memory.[33] Substituting $a_t \to \frac{1}{n_{\text{eff}}}$ and $n_t \to n_{\text{eff}}$ into the updating recurrence (B.8), we obtain the asymptotic updating equations (19) and (20) that we use in the main text.

B.1 Confidence intervals

Ambiguity about the estimated mean dividend m and the variance v^2 enters considerations of robust optimization via regions of confidence. The dividend mean μ has a t distribution,

[33] An alternative interpretation is to view $n_{\text{eff}} = 1/a$ as the effective number of observations. When estimating the mean of a normally distributed variable from n observations, the estimation variance of the mean equals $1/n$ times the estimated sample variance. Consequently, the constant of proportion a can be used to identify the effective number of degrees of freedom, n_{eff}.

see (B.1). In the one-period model, the number of degrees of freedom is n. In the general model with fading memory, the number of effective degrees of freedom is the constant $1/a$. Hence, we use $n = 1/a$ in what follows. Let $q_{t_{1/a}}(\alpha)$ denote the α-quantile of the student t distribution with $1/a$ degrees of freedom. Then the corresponding confidence interval for the μ estimate based on time t posteriors is

$$\mu \in \left[m_t - q_{t_{1/a}}(1-\alpha)s_t, m_t + q_{t_{1/a}}(1-\alpha)s_t\right] \text{ with probability } (1-\alpha).$$

This confidence interval changes its center when m_t is updated over time and its size with updates in $s_t^2 = av_t^2$. Simplifying notation, we use $\kappa = q_{t_{1/a}}(1-\alpha)$ in the main text and

$$\hat{\mu}_A \in \mathcal{M}_{\mu,t} = [m_t - \kappa\sqrt{a}v_t, m_t + \kappa\sqrt{a}v_t]. \tag{B.12}$$

Equation (B.2) implies that under fading memory

$$\frac{\frac{1}{a}v^2}{\sigma^2} \sim \chi^2_{1/a},$$

with $q_{\chi^2_{1/a}}(\alpha)$ denoting the α-quantile of a $\chi^2_{1/a}$ distribution, the symmetric confidence interval (symmetric with respect to the probability) of the dividend variance, σ^2, based on time t posteriors is given by

$$\sigma^2 \in \left[\frac{1}{aq_{\chi^2_{1/a}}(1-\frac{\alpha}{2})}v_t^2, \frac{1}{aq_{\chi^2_{1/a}}(\frac{\alpha}{2})}v_t^2\right] \text{ with probability } (1-\alpha). \tag{B.13}$$

This confidence interval changes over time only because the estimates v_t^2 are updated. The quantile dependent multipliers are constants

$$\underline{\ell} = \frac{1}{aq_{\chi^2_{1/a}}(1-\frac{\alpha}{2})}, \tag{B.14}$$

$$\bar{\ell} = \frac{1}{aq_{\chi^2_{1/a}}(\frac{\alpha}{2})}. \tag{B.15}$$

B.2 Robust optimization and ambiguity in the dividend variance

To handle ambiguity in both μ and σ^2 in the one-period model of Section 2, we need to extend the max-min optimization of Equation (4) accordingly, that is,

$$\max_{\theta_A} \min_{\hat{\mu}_A \in \mathcal{M}_\mu, \hat{\sigma}_A^2 \in \mathcal{M}_{\sigma^2}} \mathbb{E}\left[-\frac{1}{\gamma_A} e^{-\gamma_A \tilde{W}_A}\right], \tag{B.16}$$

subject to the budget constraint in equation (5). The confidence intervals for μ and σ are given in Appendix B.1 with the gain parameter a defined by the number of degrees of freedom n that is, $a = 1/n$. The inner minimization is equivalent to maximizing

$$\zeta = \gamma_A \theta_A (\hat{\mu}_A - p) - \frac{1}{2}\gamma^2 \theta^2 \hat{\sigma}_A^2 + \gamma W_A, \tag{B.17}$$

with respect to the ambiguous parameters $\hat{\mu}_A$ and $\hat{\sigma}_A^2$.

The derivative of ζ with respect to $\hat{\mu}_A$ is independent of the choice of $\hat{\sigma}_A$ and is given by

$$\frac{\partial \zeta}{\partial \hat{\mu}_A} = \gamma_A \theta_A \begin{cases} > 0 & \text{if } \theta_A > 0, \\ = 0 & \text{if } \theta_A = 0, \\ < 0 & \text{if } \theta_A < 0. \end{cases}$$

The minimum of ζ with respect to $\hat{\mu}_A$ is attained at

$$\hat{\mu}_A \begin{cases} = m - \frac{\kappa}{\sqrt{n}}\hat{\sigma} & \text{if } \theta_A > 0, \\ \in \mathcal{M}_\mu & \text{if } \theta_A = 0, \\ = m + \frac{\kappa}{\sqrt{n}}\hat{\sigma} & \text{if } \theta_A < 0. \end{cases}$$

Substituting the minimizing $\hat{\mu}_A$ into ζ, and taking the derivative with respect to $\hat{\sigma}_A$ we obtain

$$\frac{\partial \zeta}{\partial \hat{\sigma}_A} \begin{cases} = -\gamma_A \theta_A \frac{\kappa}{\sqrt{n}} - \gamma_A^2 \theta^2 \hat{\sigma}_A < 0 & \text{if } \theta_A > 0, \\ = 0 & \text{if } \theta_A = 0, \\ = \gamma_A \theta_A \frac{\kappa}{\sqrt{n}} - \gamma_A^2 \theta^2 \hat{\sigma}_A < 0 & \text{if } \theta_A < 0. \end{cases}$$

Hence, ζ is always minimized by choosing the largest attainable value of $\hat{\sigma}$ which is $\sqrt{\bar{\ell}\frac{n+1}{n}}\, v$.

The outer maximization in equation (B.16) yields

$$\theta_A = \begin{cases} = \frac{m - \frac{\sqrt{\bar{\ell}(n+1)}}{n}\kappa v - p}{\frac{n+1}{n}\gamma_A \bar{\ell} v^2} & \text{if } m - \frac{\sqrt{\bar{\ell}(n+1)}}{n}\kappa v - p > 0, \\ = 0 & \text{otherwise }, \\ = \frac{m + \frac{\sqrt{\bar{\ell}(n+1)}}{n}\kappa v - p}{\frac{n+1}{n}\gamma_A \bar{\ell} v^2} & \text{if } m + \frac{\sqrt{\bar{\ell}(n+1)}}{n}\kappa v - p < 0, \end{cases}$$

which resembles the optimal demand in (6) with

$$\begin{aligned} \gamma_A &\rightarrow \bar{\ell}\,\gamma_A, \\ \kappa &\rightarrow \sqrt{\bar{\ell}}\,\kappa. \end{aligned}$$

Therefore, including ambiguity aversion with respect to the dividend variance has no qualitative impact on the demand of type-A agents for the risky asset and is equivalent to the demand of investors who are only averse to ambiguity in the mean but have an effectively higher coefficients of risk-aversion γ_A and ambiguity aversion κ.

B.3 Stochastic and observable variance

When variance is time-varying but known, learning about the mean happens in the same way as when variance is constant but unknown. In particular, if the agents learn with fading memory, prior knowledge about the mean is discounted before updating it with new information,[34] i.e., the standard error s_{t-1}^2 increases by a factor $1/(1-a)$, mirroring the decrease in confidence in the prior as time passes and memory gradually fades. The updating recurrences in (15) to (16) become

$$m_t = (1-\eta_t)m_{t-1} + \eta_t D_t \tag{B.18}$$

$$\frac{1}{s_t^2} = \frac{1}{\frac{1}{1-a}s_{t-1}^2} + \frac{1}{\sigma_t^2} \tag{B.19}$$

$$s_t^2 = \eta_t \sigma_t^2 \quad \text{with} \quad \eta_t = \frac{\frac{1}{1-a}s_{t-1}^2}{\frac{1}{1-a}s_{t-1}^2 + \sigma_t^2}, \tag{B.20}$$

Equation (B.19) shows that with fading memory, $a > 0$, the standard error of the mean, s_t, might increase with new observations. Whenever the information content of the new dividend $(1/\sigma_t^2)$ is low compared to the fading memory effect $(s_{t-1}^2/(1-a))$, new, noisy, information might actually increase the standard error s_t. This happens when the new observation is sufficiently noisy and/or memory fades sufficiently fast, that is, when $\sigma_t^2 > s_{t-1}^2/a$. In contrast, when $\sigma_t^2 < s_{t-1}^2/a$, the information content of the new dividend observation is sufficiently large to outweigh the effect of fading memory. Without fading memory, that is, $a \rightarrow 0$, the latter effect always dominates and, as we discuss in the main text, under

[34] For the discount factor approach see West and Harrison (2006).

stochastic but known volatility, new observations always increase the precision of the mean estimate, in contrast to the case in which variance is unknown.

C Solution algorithm

To find an equilibrium for the general model described in Section 3, we extend the algorithm proposed by Dumas and Lyasoff (2012) for preference heterogeneity and learning. We first provide a summary of the methodology and then discuss explicitly our numerical implementation.

Methodology. We compute the equilibrium through backward induction. Given the endogenous variable $\omega_t = c_{i,t}/D_t$, or equivalently, the state price $\phi_{i,t} = u'_{i,t}(c_{i,t}) = u'(\omega_t D_t)$ we solve the following system of equations:

1. First-order conditions for $t+1$-consumption:

$$u'_{i,t+1}(c_{i,t+1}) = \phi_{i,t+1}, \quad 0 \le t \le T-1, \quad i = A, B \tag{C.1}$$

   with ϕ_{t+1} denoting the Lagrangian associated to $t+1$ budget constraint (25).

2. Time $t+1$ flow budget constraint (25)

$$c_{i,t+1} + F_{i,t+1} = \theta^s_{i,t}\left(p^s_{t+1} + D_{t+1}\right) + \theta^b_{i,t} \cdot 1, \quad t = 0, \ldots, T-1, \quad i = A, B \tag{C.2}$$

   where

$$F_{i,t+1} \equiv \theta^s_{i,t+1} p^s_{t+1} + \theta^b_{i,t+1} p^b_{t+1} \tag{C.3}$$

   denotes agent i's wealth *exiting* time $t+1$.

3. Kernel conditions

$$\mathbb{E}_{A,t}\left[\frac{\phi_{A,t+1}}{\phi_{A,t}} \times (p^s_{t+1} + D_{t+1})\right] = \mathbb{E}_{B,t}\left[\frac{\phi_{B,t+1}}{\phi_{B,t}} \times (p^s_{t+1} + D_{t+1})\right] \tag{C.4}$$

$$\mathbb{E}_{A,t}\left[\frac{\phi_{A,t+1}}{\phi_{A,t}} \times 1\right] = \mathbb{E}_{B,t}\left[\frac{\phi_{B,t+1}}{\phi_{B,t}} \times 1\right] \tag{C.5}$$

The expectations $\mathbb{E}_{i,t}[\cdot]$ account for the difference in agent's preferences when forming portfolios. We model type-A agents' preference for robustness as a change in agent B's probability measure. Specifically, for agent A the distribution of dividend has mean $\hat{\mu}_{A,t} = m_t \pm \kappa v_t \sqrt{a}$ and variance $\hat{\sigma}^2_{A,t} = (1+a)\bar{\ell} v_t^2$, with m_t the posterior mean from equation (19), a the constant gain parameter, and $\kappa > 0$ and $\bar{\ell} > 1$ ambiguity aversion parameters determining the confidence intervals for the mean and variance in equations (22)–(23). For agent B, the distribution of dividend has mean $\hat{\mu}_{B,t} = m_t$ and variance $\hat{\sigma}^2_{B,t} = (1+a)v_t^2$.

4. Market clearing conditions

$$\theta^s_{A,t} + \theta^s_{B,t} = 1 \tag{C.6}$$
$$\theta^b_{A,t} + \theta^b_{B,t} = 0. \tag{C.7}$$

We model uncertainty as an event tree $(\Omega, \mathbb{F})$ with Ω the state originating from the four state variables $(D_t, m_t, v_t^2, \omega_t)_{t=0}^T$, and $\mathbb{F} \equiv \{\mathcal{F}_t\}_{t=0}^T$ a filtration. We denote by $\xi \in \mathcal{F}_t$ a generic node of the event tree at time t and by $\xi^+ \subset \mathcal{F}_{t+1}$ the set of its successor nodes. To construct an equilibrium by backward induction, for each node $\xi_t \in \mathcal{F}_t$, $t = T-1, T-2, \ldots, 1$, we need to solve the conditions (C.1)-(C.7) across all the successors nodes $\eta \in \xi^+$. At time $t = 0$, we are left with equations (C.1) and (C.6) which are the only "forward" conditions, that is,

$$u'(c_{i,0}) = \phi_{i,0}, \quad i = A, B \tag{C.8}$$
$$c_{i,0} + F_{i,0} = \bar{\theta}^s_{i,0}(p^s_0 + d_0) + \bar{\theta}^b_{i,0} p^b_0, \tag{C.9}$$

where $\bar{\theta}^s_{i,0}$ and $\bar{\theta}^b_{i,0}$ are agent i's endowment of the stock and bond, and

$$F_{i,0} = \mathbb{E}_{i,0}\left[\frac{\phi_{i,1}}{\phi_{i,0}} \times (F_{i,1} + c_{i,1})\right] \tag{C.10}$$
$$p^s_0 = \mathbb{E}_{i,0}\left[\frac{\phi_{i,1}}{\phi_{i,0}} \times (p^s_1 + D_1)\right] \tag{C.11}$$
$$p^b_0 = \mathbb{E}_{i,0}\left[\frac{\phi_{i,1}}{\phi_{i,0}} \times 1\right]. \tag{C.12}$$

The system (C.8)–(C.9) needs to be solved for $\{c_{i,0}, \phi_{i,0}\}$, $i = A, B$. Because markets are incomplete, the individual stochastic discount factors, $\phi_{i,t+1}/\phi_{i,t}$, will not be equated across agents. However, all agents must agree on the price in equilibrium. Therefore, from the kernel conditions (C.4)–(C.5) we have that the exiting wealth $F_{i,t}$ and the asset prices p_t^s and p_t^b can be written in term's of either agent's state price, that is, for $t \geq 0$,

$$F_{i,t} = \mathbb{E}_{i,t}\left[\frac{\phi_{i,1}}{\phi_{i,0}} \times (F_{i,t+1} + c_{i,t+1})\right], \quad F_{i,T} \equiv 0 \tag{C.13}$$

$$p_t^s = \mathbb{E}_{i,t}\left[\frac{\phi_{i,t+1}}{\phi_{i,t}} \times (p_{t+1}^s + D_{t+1})\right], \quad p_T^s \equiv 0 \tag{C.14}$$

$$p_0^b = \mathbb{E}_{i,t}\left[\frac{\phi_{i,t+1}}{\phi_{i,t}} \times 1\right], \quad p_T^b \equiv 0. \tag{C.15}$$

Numerical implementation. We now describe the numerical implementation of the algorithm.

1. At each time t the problem is fully described by four state variables: $(D_t, m_t, v_t^2, \omega_t)$. We discretize the state space using a $N_d \times N_m \times N_v \times N_\omega$ grid. Each point of this grid represents a decision node ξ. While the consumption share is bounded between 0 and 1, the other state variables do not have natural bounds. We construct the grids for m_t and v_t^2 by centering them around the initial belief m_0 and v_0^2. Similarly, we construct the grid for D_t by centering it around m_0. A point ξ on this four-dimensional grid represents a decision node. For each decision node ξ we identify a set of successor nodes $\eta \in \xi^+$ which we refer to as "active" nodes. We determine such active nodes η together with the corresponding subjective probabilities $q_{i,\eta}$ in order to match the preferences of both agents, see (C.4) and (C.5). Specifically, we use the dividend realizations in $\eta \in \xi^+$ together with the priors in the decision node ξ and the Bayesian updating rules to calculate the range of posteriors m_t and v_t^2 so as to match the means and variances $\hat{\mu}_i$ and $\hat{\sigma}_i^2$ that agents use in forming portfolios. As a consequence of this construction, the equilibrium can be calculated by backward induction within a (hyper)cone of the discretized state space.

2. The core of the algorithm consists of solving the equilibrium conditions (C.1)–(C.7). Denoting by $q_{i,\eta}$ the subjective probabilities for the successor nodes $\eta \in \xi^+$, we obtain the current portfolio and the investors' future state prices (or consumption) in all successor nodes $\eta \in \xi^+$ by solving the following system of equations for $i = A, B$ and

for all $\eta \in \xi^+$:

$$c_{i,t+1,\eta} + F_{i,t+1,\eta} = \theta^s_{i,t,\xi}(p^s_{t+1,\eta} + D_{t+1,\eta}) + \theta^b_{i,t,\xi}, \tag{C.16}$$

$$F_{i,t} = \sum_{\eta \in \xi^+} q_{i,\eta} \left[\frac{u'_i(c_{i,t+1,\eta})}{u'_i(c_{i,t,\xi})} (F_{i,t+1,\eta} + c_{i,t+1,\eta}) \right], \tag{C.17}$$

$$p^s_t = \sum_{\eta \in \xi^+} q_{i,\eta} \left[\beta \frac{u'_i(c_{i,t+1,\eta})}{u'_i(c_{i,t,\xi})} \left(p^s_{t+1,\eta} + D_{t+1,\eta} \right) \right], \tag{C.18}$$

$$p^b_t = \sum_{\eta \in \xi^+} q_{i,\eta} \left[\beta \frac{u'_i(c_{i,t+1,\eta})}{u'_i(c_{i,t,\xi})} 1 \right], \tag{C.19}$$

$$D_{t+1,\eta} = c_{A,t+1,\eta} + c_{B,t+1,\eta}, \tag{C.20}$$

$$1 = \theta^s_{A,t,\xi} + \theta^s_{B,t,\xi}, \tag{C.21}$$

$$0 = \theta^b_{A,t,\xi} + \theta^b_{B,t,\xi}, \tag{C.22}$$

with $F_{i,t+1,\eta}$ denoting agent i's wealth exiting time $t+1$ in the successor nodes $\eta \in \xi^+$.

3. When we reach time 0, we use the initial conditions to solve for the consumption of agents. After a dividend realizations and the subsequent update, the estimated mean and variance values for the next decision stage will generally not be located exactly on our chosen discrete grid points for both agents. Therefore, we use cubic-spline approximations to capture the non-linear relationship between decision variable and state variables. The accuracy of approximation increase with the fineness of the grid.

References

Agnew, J., P. Balduzzi, and A. Sunden, 2003, "Portfolio choice and trading in a large 401(k) plan," *American Economic Review*, 93(1), 193–215.

Ai, H., and R. Bansal, 2018, "Risk preferences and the macroeconomic announcement premium," *Econometrica*, 86(4), 1383–1430.

Anderson, E., L. P. Hansen, and T. Sargent, 2000, "Robustness, detection and the price of risk," working paper, mimeo.

Bansal, R., and A. Yaron, 2004, "Risks for the long run: A potential resolution of asset pricing puzzles," *The journal of Finance*, 59(4), 1481–1509.

Barber, B. M., X. Huang, T. Odean, and C. Schwarz, 2021, "Attention-induced trading and returns: Evidence from Robinhood users," *The Journal of Finance*.

Barber, B. M., and T. Odean, 2008, "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors," *The Review of Financial Studies*, 21(2), 785–818.

Barberis, N., 2000, "Investing for the long run when returns are predictable," *The Journal of Finance*, 55(1), 225–264.

Barrot, J.-N., R. Kaniel, and D. Sraer, 2016, "Are retail traders compensated for providing liquidity?," *Journal of Financial Economics*, 120(1), 146–168.

Berkman, H., P. D. Koch, L. Tuttle, and Y. J. Zhang, 2012, "Paying attention: Overnight returns and the hidden cost of buying at the open," *Journal of Financial and Quantitative Analysis*, 47(4), 715–741.

Bewley, T. F., 2011, "Knightian decision theory and econometric inferences," *Journal of Economic Theory*, 146(3), 1134–1147.

Bhamra, H. S., and R. Uppal, 2014, "Asset prices with heterogeneity in preferences and beliefs," *The Review of Financial Studies*, 27(2), 519–580.

Buss, A., R. Uppal, and G. Vilkov, 2021, "Dynamics of Asset Demands with Confidence Heterogeneity," working paper.

Cao, H. H., T. Wang, and H. H. Zhang, 2005, "Model uncertainty, limited market participation, and asset prices," *The Review of Financial Studies*, 18(4), 1219–1251.

Cao, Z., W. Wu, and Y. Wu, 2022, "Uncertainty and the Beta Anomaly," working paper.

Chapman, D. A., and V. Polkovnichenko, 2009, “First-order risk aversion, heterogeneity, and asset market outcomes,” *The Journal of Finance*, 64(4), 1863–1887.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer, 2016a, “Asset Pricing When ‘This Time Is Different’,” *The Review of Financial Studies*, 30(2), 505–535.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer, 2016b, “Parameter learning in general equilibrium: The asset pricing implications,” *American Economic Review*, 106(3), 664–98.

Constantinides, G. M., and D. Duffie, 1996, “Asset pricing with heterogeneous consumers,” *Journal of Political Economy*, 104(2), 219–240.

Dangl, T., and M. Halling, 2012, “Predictive regressions with time-varying coefficients,” *Journal of Financial Economics*, 106(1), 157–181.

David, A., 1997, “Fluctuating confidence in stock markets: Implications for returns and volatility,” *Journal of Financial and Quantitative Analysis*, 32(4), 427–462.

Dumas, B., 1989, “Two-person dynamic equilibrium in the capital market,” *The Review of Financial Studies*, 2(2), 157–188.

Dumas, B., A. Kurshev, and R. Uppal, 2009, “Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility,” *The Journal of Finance*, 64(2), 579–629.

Dumas, B., and A. Lyasoff, 2012, “Incomplete-Market Equilibria Solved Recursively on an Event Tree,” *The Journal of Finance*, 67(5), 1897–1941.

Dunn, O. J., 1964, “Multiple comparisons using rank sums,” *Technometrics*, 6(3), 241–252.

Easley, D., and M. O’Hara, 2009, “Ambiguity and Nonparticipation: The Role of Regulation,” *The Review of Financial Studies*, 22(5), 1817–1843.

Epstein, L. G., and M. Schneider, 2008, “Ambiguity, information quality, and asset pricing,” *The Journal of Finance*, 63(1), 197–228.

Evans, G. W., and S. Honkapohja, 2009, “Learning and macroeconomics,” *Annual Review of Economics*, 1(1), 421–449.

——— , 2012, *Learning and expectations in macroeconomics*. Princeton University Press.

Fama, E. F., and J. D. MacBeth, 1973, “Risk, return, and equilibrium: Empirical tests,” *The Journal of Political Economy*, 81(3), 607–636.

Fox, C. R., and A. Tversky, 1995, “Ambiguity aversion and comparative ignorance,” *The Quarterly Journal of Economics*, 110(3), 585–603.

Frazzini, A., and O. A. Lamont, 2007, "The earnings announcement premium and trading volume," *NBER working paper*.

Garlappi, L., R. Uppal, and T. Wang, 2007, "Portfolio selection with parameter and model uncertainty: A multi-prior approach," *The Review of Financial Studies*, 20(1), 41–81.

Gârleanu, N., and S. Panageas, 2015, "Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing," *Journal of Political Economy*, 123(3), 670–685.

Gilboa, I., and D. Schmeidler, 1989, "Maxmin Expected Utility Theory with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141—153.

Glossner, S., P. Matos, S. Ramelli, and A. F. Wagner, 2020, "Where do institutional investors seek shelter when disaster strikes? Evidence from COVID-19," *SSRN working paper*.

Graham, J. R., C. R. Harvey, and H. Huang, 2009, "Investor competence, trading frequency, and home bias," *Management Science*, 55(7), 1094–1106.

Greene, W. H., 2020, *Econometric Analysis*. Pearson, 8 edn.

Hansen, L., and T. J. Sargent, 2001, "Robust control and model uncertainty," *American Economic Review*, 91(2), 60–66.

Hansen, L. P., and T. J. Sargent, 2008, *Robustness*. Princeton university press.

Heath, C., and A. Tversky, 1991, "Preference and belief: Ambiguity and competence in choice under uncertainty," *Journal of Risk and Uncertainty*, 4(1), 5–28.

Hirshleifer, D. A., J. N. Myers, L. A. Myers, and S. H. Teoh, 2008, "Do individual investors cause post-earnings announcement drift? Direct evidence from personal trades," *The Accounting Review*, 83(6), 1521–1550.

Hu, G., R. D. McLean, J. Pontiff, and Q. Wang, 2014, "The year-end trading activities of institutional investors: Evidence from daily trades," *The Review of Financial Studies*, 27(5), 1593–1614.

Illeditsch, P. K., J. V. Ganguli, and S. Condie, 2021, "Information inertia," *The Journal of Finance*, 76(1), 443–479.

Ilut, C. L., and M. Schneider, 2022, "Modeling Uncertainty as Ambiguity," *NBER working paper*.

Kaniel, R., S. Liu, G. Saar, and S. Titman, 2012, "Individual Investor Trading and Return Patterns around Earnings Announcements," *The Journal of Finance*, 67(2), 639–680.

Kaniel, R., G. Saar, and S. Titman, 2008, "Individual investor trading and stock returns," *The Journal of Finance*, 63(1), 273–310.

Koop, G., and D. Korobilis, 2013, "Large time-varying parameter VARs," *Journal of Econometrics*, 177, 185–198.

Krusell, P., and A. A. Smith, Jr, 1998, "Income and wealth heterogeneity in the macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

Kruskal, W. H., and W. A. Wallis, 1952, "Use of ranks in one-criterion variance analysis," *Journal of the American Statistical Association*, 47(260), 583–621.

Leippold, M., F. Trojani, and P. Vanini, 2008, "Learning and asset prices under ambiguous information," *The Review of Financial Studies*, 21(6), 2565–2597.

Lewellen, J., and J. Shanken, 2002, "Learning, asset-pricing tests, and market efficiency," *The Journal of Finance*, 57(3), 1113–1145.

Li, C. W., A. Tiwari, and L. Tong, 2017, "Investment decisions under ambiguity: Evidence from mutual fund investor behavior," *Management Science*, 63(8), 2509–2528.

Malmendier, U., and S. Nagel, 2011, "Depression babies: do macroeconomic experiences affect risk taking?," *The Quarterly Journal of Economics*, 126(1), 373–416.

Mankiw, N. G., 1986, "The equity premium and the concentration of aggregate shocks," *Journal of Financial Economics*, 17(1), 211–219.

Nagel, S., 2012, "Evaporating liquidity," *The Review of Financial Studies*, 25(7), 2005–2039.

Nagel, S., and Z. Xu, 2021, "Asset Pricing with Fading Memory," *The Review of Financial Studies*.

——— , 2022, "Dynamics of Subjective Risk Premia," .

Panageas, S., 2020, "The implications of heterogeneity and inequality for asset pricing," working paper, National Bureau of Economic Research.

Pástor, L., 2000, "Portfolio selection and asset pricing models," *The Journal of Finance*, 55(1), 179–223.

Pástor, L., and P. Veronesi, 2009, "Learning in Financial Markets," *Annual Review of Financial Economics*, 1, 361–381.

Pástor, L., and M. B. Vorsatz, 2020, "Mutual fund performance and flows during the COVID-19 crisis," *The Review of Asset Pricing Studies*, 10(4), 791–833.

Savor, P., and M. Wilson, 2016, "Earnings announcements and systematic risk," *The Journal of Finance*, 71(1), 83–138.

Segal, U., and A. Spivak, 1990, "First order versus second order risk aversion," *Journal of Economic Theory*, 51(1), 111–125.

Stambaugh, R. F., 2014, "Presidential address: Investment noise and trends," *The Journal of Finance*, 69(4), 1415–1453.

Veronesi, P., 1999, "Stock Market Overreaction on Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12(5), 975–1007.

West, M., and J. Harrison, 2006, *Bayesian Forecasting and Dynamic Models*. Springer Science & Business Media.

Xia, Y., 2001, "Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation," *The Journal of Finance*, 56(1), 205–246.